

Matching Auctions*

Daniel Fershtman[†] Alessandro Pavan[‡]

February 2019

Abstract

This paper is about (platform) mediated matching in markets in which valuations evolve over time. We introduce and then study a class of dynamic matching auctions where, in each period, agents from two sides of a market submit multiple bids, one for each possible partner. Each match receives a “score” that is a weighted average of the involved agents’ reciprocal bids, net of the platform’s match-specific costs. The weights are determined by the agents’ membership statuses and vary with the platform’s objectives. In each period, the matches that maximize the sum of the bilateral scores subject to individual and aggregate capacity constraints are implemented. We show that, under appropriate conditions, this class includes both welfare- and profit-maximizing mechanisms. When match values are positive and none of the capacity constraints binds, profit maximization results in fewer interactions than welfare maximization, in each period. This conclusion need not extend to markets in which individual and/or aggregate capacity constraints bind and/or agents dislike certain interactions. Finally, we discuss how similar auctions but with forward-looking “index scores” can be used in markets where match values depend on past interactions, for example due to experimentation, a preference for variety, or habit formation.

JEL Classification Numbers: D82, C73, L1.

Keywords: matching, platforms, asymmetric information, learning, dynamic scoring auctions.

1 Introduction

Motivation. In the last few years, matching markets have been growing at an unprecedented rate, reflecting the role that the “sharing economy” is taking in the organization of modern

*For helpful comments and suggestions, we thank Dirk Bergemann, Yeon-Koo Che, Eddie Dekel, Drew Fudenberg, Benjamin Hermalin, Terence Johnson, Fuhito Kojima, Benny Moldovanu, Larry Samuelson, Rakesh Vohra, Jean Tirole, Asher Wolinsky, Leeat Yariv, and seminar participants at various universities and conferences where the paper was presented. Matteo Camboni provided excellent research assistance. Pavan also thanks the National Science Foundation for financial support under the grant SES 1530798 and Bocconi University for its hospitality during the 2017-2018 academic year. The usual disclaimer applies.

[†]Tel-Aviv University and University of Bonn (dfershtm@uni-bonn.de)

[‡]Northwestern University (alepavan@northwestern.edu)

business activities. In e-commerce, for example, a sizable fraction of trade is mediated by business-to-business (B2B) platforms matching vendors with procurers in search of business opportunities. Likewise, a sizable fraction of online advertising is mediated by ad exchanges, search engines, media outlets, online malls, and videogame consoles, matching advertisers with either consumers or content providers.¹

Platform mediated matching plays an important role also in the growing market for scientific outsourcing where intermediaries such as Science Exchange match labs with idle equipment with research units wishing to conduct experiments off-site;² project finance, where consulting firms match startups with lenders;³ lobbying, where commercial firms match interest groups with policy makers;⁴ the market for private medical-tourism services, where intermediaries such as MEDIGO match patients from abroad with US physicians providing specialized treatments;⁵ and the market for organized events, where online platforms such as meetings.com match clients in search of conference venues, meeting spaces, corporate travel plans, or other hospitality services, with providers supplying such services.⁶

Matching in these markets is typically many-to-many and agents change partners over time. In such dynamic matching environments, how do allocations and pricing respond to variations in preferences and/or the arrival of information? And how do the matching services provided by private (profit-maximizing) intermediaries reflect market power and compare to the matches that maximize welfare? These questions are receiving a lot of attention from market designers, analysts, and policy makers alike. Many predict that matching intermediaries will eventually resort to auctions to mediate the interactions between the various sides of a market. In fact, auctions are already used by online ad exchanges such as Google's DoubleClick and Microsoft's Exchange to match advertisers with content providers. In such auctions, advertisers bid repeatedly over time to place their ads on the website of multiple content providers and, over time, content providers ask different ad-specific prices to display the ads.⁷ Such auctions often involve the use of dynamic "compatibility scores" whereby the bids of agents from different

¹See, e.g., "Marketing in the digital age: A brand new game", *The Economist*, August 29, 2015.

²See, e.g., "Uber for Experiments," *The Economist*, December 6, 2014.

³A similar role is played by peer-to-peer lending. Platforms such as Prosper and LendingClub match borrowers with individual or institutional lenders. Prosper operated as an online auction marketplace between 2006 and 2009. It recently switched to a system of pre-set rates determined by an algorithm that accounts for the borrowers' credit risk.

⁴See Allard (2008) and Kang and You (2016) for how lobbying firms provide tailored (many-to-many) matching services and dynamically price-discriminate each side of the lobbying market. See also Dekel, Jackson and Wolinsky (2008) for a detailed account of how intermediaries help to buy and sell votes.

⁵See <http://www.nytimes.com/2013/08/07/us/the-growing-popularity-of-having-surgery-overseas.html>.

⁶As in the case of most other platforms, meetings.com does not simply match the parties; it also offers tailored services, such as site selection, contract negotiations, and on-site event management.

⁷See, e.g., Mansour, Muthukrishnan, and Nisan (2012).

sides of the market are weighted so as to optimize the fit and efficacy of the advertisement without damaging broadcaster reputation.⁸ Finally, the participating agents from both sides of the market are typically charged fees to join the platform. Auctions are also used in the market for personalized display ads by search engines such as Yahoo! and Google. These intermediaries have originally used variations of the second-price auction, the so-called GSP auction (Generalized Second Price). In 2012, however, Google switched from the GSP auction to a VCG auction on the grounds that dynamic bid re-optimization is easier under the VCG protocol.⁹

The purpose of this paper is to investigate the key trade-offs that platforms face in the design of their matching protocols in markets in which match values evolve over time, as the result of learning about unknown match quality, or preferences dynamics. We view the contribution as threefold. First, we investigate the equilibrium properties of a class of dynamic matching auctions that are fairly simple to operate and accommodate for a flexible specification of the “scores” the platforms may use to weight the bids of the participating agents, as a function of their objectives. Second, we show that such a class includes auctions that maximize the platform’s profits, as well as auctions that maximize welfare. Third, we investigate how market power affects matching dynamics by comparing the matches that are sustained under profit maximization to their counterparts under welfare maximization.

Model Ingredients. The key ingredients of the model are the following. The (flow) payoff that each agent derives from each possible match is governed by two components: a time-invariant vertical characteristic that is responsible for the overall importance the agent assigns to interacting with partners from the opposite side of the market (the “vertical” type); and a vector of time-varying relation-specific values capturing the evolution of the agent’s information and preferences for specific partners (the “horizontal” types). The latter values evolve stochastically over time and may turn negative, reflecting the idea that agents may dislike certain interactions. Both the vertical and the horizontal types are the agents’ private information. The agents learn the vertical types prior to joining the platform and learn the horizontal types over time.

The model allows for limits on both the number of matches that each agent may take within each period (individual capacity constraints), as well as limits on the total number of matches that the platform can accommodate within each period (aggregate capacity constraint). Such limits may reflect time, resource, or facility constraints, but also capture certain non-separabilities in the agents’ preferences.

⁸See, e.g., https://support.google.com/adxseller/answer/2913506?hl=en&ref_topic=3376095.

⁹See, e.g., Edelman, Ostrovsky, and Schwarz (2007), Gomes and Sweeney (2014), Harris and Varian (2014), and Arnosti, Beck and Milgrom (2016).

Results. The matching auctions we consider work as follows. Upon joining the platform, each agent is invited to select a membership status whose level determines the weight the agent’s bids receive in the auctions. At any subsequent period, each agent is then asked to submit a vector of bids, one for each possible partner from the opposite side of the market (counterpart). Each bilateral match then receives a “score” that depends on the involved agents’ reciprocal bids and membership status. The score assigned to each match coincides with the weighted sum of the agents’ bids, with the weights taking different forms depending on the platform’s objective, net of possible costs to the platform of implementing the match. We allow the platform’s costs to take on negative values reflecting the possibility that the platform may benefit from certain interactions, for example because they may help advertise the platform’s matching capabilities. In each period, the platform then implements the matches that maximize the sum of the bilateral scores, taking individual and aggregate capacity constraints into account. As in the VCG and in the GSP auction, the payments the platform asks of each agent reflect the externalities the agent imposes on others due to the individual and aggregate capacity constraints. Contrary to the VCG and the GSP auction, however, such externalities may also account for the effects of matching on the agents’ informational rents. In particular, the platform may find it optimal to subsidize certain interactions and favor matches that generate lower surplus when this permits the platform to raise more revenue. In addition to charging the agents for the matches they receive over time, the platform also charges each agent a participation fee at the time of joining that depends on the selected membership status. The pricing of such membership fees takes into consideration the intertemporal surplus expected by each agent at the time of joining, adjusted by a discount that relates to the informational rent the platform leaves to the agent to induce self selection.

Our first result shows that in the proposed auctions, at all histories, including those off-path, equilibrium bidding is truthful. Bidding the myopic values is optimal for the agents because the matches under truthful bidding maximize a weighted sum of all agents’ current and future payoffs, net of the platform’s matching costs and, in case of profit-maximization, net of agents’ information rents. As in the VCG auction, this property, together with the fact that the payments make agents’ payoffs *proportional* to their contribution to total weighted surplus, guarantees that each agent finds it optimal to stay in the mechanism and bid truthfully at all periods, irrespective of the agent’s beliefs about other agents’ current and past types, and independently of past matches. As a result, the proposed auctions can be made fully transparent: At the end of each period, all membership statuses and bids are disclosed.

That, in the proposed auctions, agents find it optimal to bid truthfully at each period is fairly straightforward and follows from arguments similar to those in the literature on VCG

mechanisms. That agents have incentives to participate and select the right membership status is less obvious. We show that this follows from the combination of two monotonicities. First, under truthful bidding, the net present value of the match quality that each agent expects upon joining the platform and selecting the membership status designed for his true vertical type is nondecreasing in the agent’s vertical type. This property guarantees that if low types find it optimal to participate, so do higher ones. Second, when all agents bid truthfully in all periods, each agent’s expected match quality at the time of joining is nondecreasing in the agent’s selected status, for any profile of true vertical types. In a matching environment, such monotonicities should not be taken for granted. In fact, contrary to standard screening problems, each agent plays the double role of a buyer and of an input provider for the matches the platform sells to the other side. Furthermore, the private information each agent receives in each period is multidimensional, and while certain dimensions may contribute to higher match quality, others may contribute to a lower one. Lastly, agents may dislike certain interactions (i.e., experience a payoff below their outside option) and, for such interactions, a larger vertical type may imply a larger loss.¹⁰ Notwithstanding such complications, we show that, in the proposed auctions, a higher status brings higher match quality, both when a higher status corresponds to a higher true vertical type and when it does not. Both monotonicities are key to our results.

As mentioned above, the proposed auctions are fairly simple and allow for different specifications of the weights the platforms may use to aggregate the bids of agents from different sides of the market when computing the bilateral scores.¹¹ Our second and third results show that, under reasonable conditions, such class includes auctions that maximize the platform’s profits, as well as auctions that maximize welfare, *over all possible mechanisms*. We then use these results to shed light on the inefficiencies brought in by market power in dynamic matching markets. We show that, in markets in which all agents assign a nonnegative value to all interactions and none of the capacity constraints binds, in each period, each match active under profit maximization is also active under welfare maximization. In the presence of capacity constraints, instead, certain matches active under profit maximization need not be active under welfare maximization. However, when the only binding capacity constraint is the aggregate one (i.e., the platform’s), the total number of active matches in each period under profit maximization is always inefficiently low. Interestingly, this property does not extend to markets where individual capacity constraints may be binding. In this case,

¹⁰Technically, the agents’ payoffs at the time of joining need not satisfy the increasing-difference property between the agents’ vertical types and the sequence of the matches.

¹¹The weights in our auctions play a role similar to that played by the “compatibility scores” that ad exchanges use to select the matches between advertisers and content providers (see, e.g., Chapter 6 in Moghaddam and Shimon (2016)).

profit maximization may result in an inefficiently large number of matches, for any number of periods. The same is true in markets in which agents dislike certain interactions (that is, derive a payoff lower than their outside option from certain matches). The last few years have witnessed great interest (from policy makers and academics alike) on how to regulate matching intermediaries.¹² We believe the above results may provide some guidance on how to intervene in such markets.

The analysis in the baseline model is conducted under the assumption that match values evolve over time in response to the exogenous arrival of information and/or exogenous preferences shocks. At the end of the paper, however, we discuss how similar auctions can be used in certain markets where the evolution of the match values is endogenous, be it the result of experimentation (whereby agents learn the attractiveness of their partners by interacting with them), a preference for variety (whereby agents gradually lose interest in those partners they already interacted with), or habit formation (whereby match values increase with the number of past interactions). In these markets, the bilateral scores take the form of forward-looking “indexes” that account for the benefit of generating new information (in the case of experimentation), the opportunity cost of reducing future match values (in the case of a preference for variety), or the value of enhancing future match values (in the case of habit formation).

With endogenous processes, the conditions under which our matching auctions are optimal (i.e., profit- or welfare-maximizing) are restrictive. In particular, they require the match values to evolve independently across matches, each match value to remain frozen when the match is not active, and the capacity constraints to be either lax or imposing that a single match be formed in each period.¹³ Such conditions are typical in dynamic problems with endogenous processes. Notwithstanding these limitations, we expect the insights to be useful in a broader class of dynamic markets with endogenous match values.

Outline. The rest of the paper is organized as follows. Section 2 describes the environment. Section 3 introduces our dynamic matching auctions. Section 4 derives equilibrium properties of the proposed mechanisms. Section 5 identifies a subclass of our matching auctions that maximize the platform’s profits and a class that maximize welfare, and contrasts matching dynamics under profit maximization with their counterparts under welfare maximization. Section 6 discusses the case of endogenous processes. Section 7 reviews the pertinent literature and further discusses the paper’s contribution vis-a-vis previous work, while Section 8 concludes. Proofs omitted in the main text are in the Appendix at the end of the document.

¹²See, e.g., “Online Platforms: Nostrums for Rostrums” and “Regulating Technology Companies: Taming the Beasts,” *The Economist*, May 28, 2016.

¹³In a previous version, we discussed how such restrictions can be relaxed a little bit, albeit at the cost of an increase in complexity.

2 The Environment

Agents, Matches, and Preferences. A platform mediates the interactions among agents from two sides of a market, A and B . There are $n^A \in \mathbb{N}$ agents on side A and $n^B \in \mathbb{N}$ agents on side B , with $N^A \equiv \{1, \dots, n^A\}$ and $N^B \equiv \{1, \dots, n^B\}$ denoting the corresponding sets of agents on the two sides. Time is discrete, indexed by $t = 0, 1, \dots, \infty$. Agents live for infinitely many periods and can change partners infinitely many times.

To ease the exposition, below we describe various features of the environment by focusing on a generic agent from side A . It should be understood that a similar description applies to the side- B agents.

The flow period- t payoff that agent $i \in N^A$ derives from being matched to agent $j \in N^B$ is given by

$$v_{ijt}^A(\theta_i^A, \varepsilon_{ijt}^A) = \theta_i^A \cdot \varepsilon_{ijt}^A. \quad (1)$$

The parameter θ_i^A is time- and match-invariant and captures the overall importance that agent i assigns to interacting with agents from the opposite side of the market. That is, θ_i^A parametrizes the value that agent i assigns to interacting with a generic agent from side B , prior to conditioning on the specific profile of the latter agent. The parameter ε_{ijt}^A , instead, is match-specific and time-variant and captures the attractiveness of agent j in the eyes of agent i . These match-specific values evolve over time, reflecting the change in the agents' true, or perceived, attractiveness. They can either represent the evolution of the agents' beliefs about fixed, but unknown, match qualities, or variations in attractiveness triggered by stochastic changes in the environment. Hereafter we refer to θ_i^A as the agent's “vertical type” and to $\varepsilon_{it}^A \equiv (\varepsilon_{ijt}^A)_{j \in N^B}$ as the profile of the agent's period- t “horizontal types”. We refer to $v_{it}^A \equiv (v_{ijt}^A)_{j \in N^B}$ as the agent's period- t “match values”.

Both the agent's vertical and horizontal types are his own private information. The vertical type θ_i^A is drawn from an absolutely continuous cumulative distribution function F_i^A with density $f_i^A(\theta_i^A) > 0$ if, and only if, $\theta_i^A \in \Theta_i^A \equiv [\underline{\theta}_i^A, \bar{\theta}_i^A]$, with $\underline{\theta}_i^A > 0$. Vertical types are drawn independently across agents and from the horizontal types $\varepsilon \equiv (\varepsilon_{ijt}^k)_{t=1, \dots, \infty}^{(i,j) \in N^A \times N^B, k=A,B} \in \mathcal{E}$. Importantly, while we restrict the agents' vertical types to be nonnegative, we allow the horizontal types to be negative, reflecting the possibility that agents may dislike certain interactions (that is, may derive a payoff lower than their outside option from interacting with such agents).¹⁴

Agents learn their vertical types prior to joining the platform, whereas they learn their

¹⁴Under the assumed multiplicative structure $v_{ijt}^A = \theta_i^A \cdot \varepsilon_{ijt}^A$, allowing the vertical types to also take on negative values would introduce confusion given that the horizontal types ε_{ijt}^A are already allowed to take on negative values.

horizontal types gradually over time, with each ε_{ijt}^A learned at the beginning of period t , $t = 1, \dots, \infty$. This assumption is motivated by the idea that, in many markets of interest, agents learn their preferences for specific partners only after “getting on board”.

For any $t \geq 1$, and any pair of agents $(i, j) \in N^A \times N^B$, let $X_{ijt} \equiv \{0, 1\}$, with $x_{ijt} = 1$ in case the pair is matched in period t , and with $x_{ijt} = 0$ in case the pair is unmatched.

All agents have expected-utility preferences and maximize the expected discounted sum of their flow payoffs using the common discount factor $\delta \in (0, 1]$. Let $p_t \equiv (p_{it}^k)_{k=A,B}$ denote the payments collected by the platform from the two sides of the market in period t , and $p \equiv (p_t)_{t=0}^\infty$ an entire sequence of payments. Note that, while the matching starts in period one, the platform may start collecting payments from the agents in period zero, when the agents join the platform. Also note that payments are allowed to be negative, reflecting the possibility that (a) agents may dislike certain interactions and ask to be compensated, or (b) even if all agents like interacting with all other agents, the platform may want to cross-subsidize certain interactions.

Agent i 's (Bernoulli) payoff function is given by

$$U_i^A = \sum_{t=1}^{\infty} \delta^t \sum_{j \in N^B} x_{ijt} v_{ijt}^A - \sum_{t=0}^{\infty} \delta^t p_{it}^A. \quad (2)$$

The platform's (Bernoulli) payoff function is given by

$$U_0 = \sum_{t=0}^{\infty} \delta^t \left(\sum_{i \in N^A} p_{it}^A + \sum_{j \in N^B} p_{jt}^B \right) - \sum_{t=1}^{\infty} \delta^t \left(\sum_{i \in N^A} \sum_{j \in N^B} c_{ijt} \cdot x_{ijt} \right).$$

The platform's payoff is thus equal to the discounted sum of the payments collected from the two sides of the market, net of possible costs of implementing the matches, where $c_{ijt} \in \mathbb{R}$ is the period- t "cost" of matching the pair (i, j) . As anticipated in the Introduction, we allow the costs to take on negative values so as to capture the possibility that the platform may derive positive benefit from certain matches, for example because they can be used to advertise the platform's matching capabilities. These costs also incorporate all the auxiliary services that the platform may provide to the two sides, over and above mediating the matching.¹⁵

Evolution of Match Values. The match values $v_i^A \equiv (v_{ijt}^A)_{j \in N^B}$ that agent i derives from possible partners from side B are correlated over time, both through the fully persistent vertical type θ_i^A and through the partially persistent horizontal types $\varepsilon_{it}^A \equiv (\varepsilon_{ijt}^A)_{j \in N^B}$. Each vector ε_{it}^A is drawn from $\mathcal{E}_{it}^A \equiv \prod_{j=1}^{n^B} \mathcal{E}_{ijt}^A$ according to the cdf $G_{it}^A(\varepsilon_{it}^A | \varepsilon_{it-1}^A)$. While not essential

¹⁵The case where the costs depend on past matches is discussed in Section 6.

to the results, we find it convenient to think of each $\mathcal{E}_{ijt}^A \subseteq \mathbb{R}$ either as the entire real line, or as a compact and connected subset of it. Importantly, while the support of G_{it}^A is a subset of \mathcal{E}_{it}^A , we allow for the possibility that, for certain histories ε_{it-1}^A , it is a strict subset of \mathcal{E}_{it}^A , with $\mathcal{E}_{i0}^A \equiv \emptyset$.

We let $G_i^A \equiv (G_{it}^A)_{t=1}^\infty$ and then denote by $G \equiv (G_i^k)_{i=1, \dots, n^k}^{k=A, B}$ the complete collection of kernels for all agents. To guarantee that each agent's expected payoff is well defined at all histories and satisfies a certain envelope formula (more below), we assume that, for all $i \in N^A$, $(G_{it}^A)_{t=1}^\infty$ is such that $\mathbb{E} \left[\sum_{t=1}^\infty \delta^t \sum_{j \in N^B} |\varepsilon_{ijt}^A| \right] \leq E_i^A$ for some constant $E_i^A > 0$.

Capacity Constraints. At each period $t \geq 1$, agent i can be matched to at most m_i^A agents from side B . The special case of a market where matching is *one-to-one* corresponds to a market in which $m_i^k = 1$ all $i = 1, \dots, n^k$, $k = A, B$, whereas the special case of a market where matching is *many-to-many* with no binding individual capacity constraints corresponds to a market in which $m_i^k \geq n^{-k}$, all $i = 1, \dots, n^k$, $k = A, B$.¹⁶ In addition, in each period, the platform can match at most M pairs of agents. The platform can delete some of the previously formed matches and create new ones. The total number of existing matches, however, cannot exceed M in each period. For example, the number of ads and articles jointly displayed in each period on a given media outlet may be naturally limited by the outlet's physical capacity. More generally, the above limits may reflect space, time, and resource constraints, but also capture certain non-separabilities in payoffs (in the case of individual constraints).

A period- t match $x_t \in \prod_{(i,j) \in N^A \times N^B} X_{ijt}$ is *feasible* if (1) for each $i = 1, \dots, n^A$, $\sum_{j \in N^B} x_{ijt} \leq m_i^A$, (2) for each $j = 1, \dots, n^B$, $\sum_{i \in N^A} x_{ijt} \leq m_j^B$, and (3) $\sum_{i \in N^A} \sum_{j \in N^B} x_{ijt} \leq M$. We denote by X_t the set of feasible period- t matches, and by $X \equiv \prod_{t=1}^\infty X_t$ the set of sequences of feasible matching allocations.

Remarks. The independence assumption about the values that different agents from the same side derive from interacting with the same counterpart avoids the possibility that the platform trivially extracts the entire surplus from each agent using payments à la Cremer and McLean (1988).

Similarly, the assumption that agents possess some private information, θ_i^A , prior to joining the platform guarantees that the matching dynamics under profit maximization differ from their counterparts under welfare maximization. In environments in which the platform maximizes welfare (a possibility we consider below), the information the agents possess at the time they join the platform plays no role (they can be assumed to know only their vertical types θ , or also the period-1 horizontal types ε_1 , or to have no private information at all). If, instead,

¹⁶Hereafter, the notation $-k$ is used to denote the side opposite to side k . So, for $k = A$, $-k = B$, whereas for $k = B$, $-k = A$.

the platform maximizes profits, then the private information the agents possess at the time they join plays a fundamental role. If the agents possess no private information at all, the platform can extract the entire surplus. If, in addition to their vertical types, they also possess private information about their period-1 horizontal types ε_1 , the optimal mechanism is significantly more complicated, because of the multi-dimensionality of the agents' initial private information (for example, it may be impossible to determine which participation constraints bind). To maintain tractability, we thus assume that the agents' private information at the time of joining is unidimensional and coincides with the vertical types θ ; however, we allow the private information the agents receive in the subsequent periods to be multi-dimensional.

The assumption that the value agent i derives from interacting with agent j is invariant to the composition of agent i 's matching set (that is, it does not depend on who else the agent interacts with) favors a certain simplicity in the definition of the scoring rules. The individual capacity constraints already introduce certain non-separabilities in the agents' payoffs. The analysis can be extended to accommodate for richer non-separabilities, albeit at the cost of an increase in the complexity of the scores.

Finally, the assumption that the horizontal types are drawn from a Markov process, as opposed to a more general process, simplifies the notation, but is irrelevant for the results.

We see the combination of the above assumptions as a convenient way to retain tractability in a fairly rich and complex environment while at the same time permitting us to shed light on the key trade-offs that platforms face in the design of their dynamic matching protocols.

Matching Mechanisms. As anticipated above, we are interested in the matches that can be sustained with simple dynamic auctions where bidders repeatedly bid for all possible partners and, in each period, the matches with the highest bilateral score are implemented. We find such auctions intuitive and appealing for the type of markets they are meant for. However, to establish the optimality of a specific subset of such auctions, we will need to compare them to arbitrary matching mechanisms, which we conveniently define below. A *matching mechanism* $\Gamma \equiv (\mathcal{M}, \mathcal{S}, \chi, \psi, \rho)$ consists of: (i) a collection of message sets $\mathcal{M} \equiv (\mathcal{M}_t)_{t=0}^\infty$, with each $\mathcal{M}_t \equiv \prod_{l \in N^k, k=A,B} \mathcal{M}_{lt}^k$ defining the set of messages the agents can send in period t ; (ii) a collection of sets of signals $\mathcal{S} \equiv (\mathcal{S}_t)_{t=0}^\infty$ that the platform may disclose to the agents, with $\mathcal{S}_t \equiv \prod_{l \in N^k, k=A,B} \mathcal{S}_{lt}^k$; (iii) a *matching rule* $\chi \equiv (\chi_t)_{t=1}^\infty$ describing, for each $t \geq 1$, the matches $\chi_t(m^t)$ implemented given the history of received messages, with $\mathcal{M}^t \equiv \prod_{s=0}^t \mathcal{M}_s$ and $\chi_t : \mathcal{M}^t \rightarrow X_t$; (iv) a *payment rule* $\psi \equiv (\psi_t)_{t=0}^\infty$ describing, for each $t \geq 0$, the payments $\psi_t(m^t)$ asked from, or made to, the agents, with $\psi_t : \mathcal{M}^t \rightarrow \mathbb{R}^{N^A+N^B}$; and (v) a *disclosure policy* $\rho \equiv (\rho_t)_{t=0}^\infty$ specifying the information $\rho(m^t)$ disclosed to the agents over time, with $\rho_t : \mathcal{M}^t \rightarrow \mathcal{S}_t$. Each ρ_t must reveal to each agent his own matches and payments. It may also

reveal additional information, but it cannot conceal the matches the individual is involved in, or the payments from/to the individual.

Each agent chooses whether or not to join the mechanism after observing his vertical type θ_l^k but before observing his horizontal types. Upon joining the mechanism, at each period $t \geq 0$, each agent $i \in N^A$, after learning $(\varepsilon_{ijt}^A)_{j \in N^B}$, sends a message m_{it}^A from the set \mathcal{M}_{it}^A (similarly for each agent $j \in N^B$, from side B). In the auctions we introduce in the next section, such messages correspond to the selection of a membership status along with a collection of bids, one for each possible partner from the opposite side.

A matching mechanism Γ is *feasible* if, for any sequence of messages $m \in \mathcal{M}$, the implemented allocations are feasible, that is, $\chi(m) \in X$.¹⁷

Solution concept. We use perfect Bayesian equilibrium (PBE) as our solution concept. The equilibria in the proposed matching auctions in the next section will satisfy properties stronger than those required by PBE. In particular, the equilibrium strategies remain optimal no matter the information each agent may possess about other agents' past messages and matching allocations, and no matter the beliefs the agent may have about other agents' past and current types. Consistently with the rest of the dynamic mechanism design literature, we will refer to such equilibria as *periodic ex-post* (see, e.g., Bergemann and Valimaki (2010), Athey and Segal (2013), and Pavan, Segal, and Toikka (2014)). Lastly, while we only require that a mechanism induces participation in period zero, in the specific auctions we propose in the next section, agents find it optimal to remain in the mechanism after each history.

3 Matching Auctions

We now introduce a class of matching mechanisms in which (a) upon joining the platform, agents are invited to select a membership status, (b) in each subsequent period, each agent is asked to submit a vector of bids, one for each possible partner, (c) bids are aggregated into a collection of *bilateral scores*, one for each match, with each score depending only on the involved agents' reciprocal current bids and membership status, (d) in each period, the matches maximizing the sum of the bilateral scores subject to individual and aggregate capacity constraints are implemented.

Definition 1 (Matching Auctions). In a *matching auction*:

- At $t = 0$, each agent $l \in N^k$, from each side $k = A, B$, is asked to select a membership status. Each status is conveniently indexed by the vertical type $\theta_l^k \in \Theta_l^k$ it is designed

¹⁷Note that the rule χ is deterministic. This is because, in this environment, the platform never gains from inducing random matches — see the discussion in the proof of Theorem 2 below.

for. An agent’s status determines the weight the platform assigns to the agent’s bids in the subsequent auctions. The payment $p_{l_0}^k$ each agent is asked to make to the platform upon joining depends on the selected membership status, according to the rule $\psi_{l_0}^k(\cdot)$ described below.

- In each subsequent period $t \geq 1$, each agent $l \in N^k$, from each side $k = A, B$, is then asked to submit a vector of bids $b_{lt}^k \in \mathbb{R}^{n-k}$, one for each possible partner. Each pair of agents $(i, j) \in N^A \times N^B$ is then assigned a “score”

$$S_{ijt}^\beta \equiv \beta_i^A(\theta_i^A) \cdot b_{ijt}^A + \beta_j^B(\theta_j^B) \cdot b_{ijt}^B - c_{ijt}, \quad (3)$$

where $\beta \equiv (\beta_i^k(\cdot))_{i \in N^k, k=A, B}$ are non-decreasing, strictly positive, and bounded functions describing the weights the platform assigns to the agents’ bids as a function of their membership status. The auction then implements the matches maximizing the sum of the bilateral scores subject to individual and aggregate capacity constraints with ties broken arbitrarily.¹⁸ All agents who are unmatched in period t , pay nothing. Matched agents make (or receive) payments $p_{lt}^k \in \mathbb{R}$ according to the rule $\psi_{lt}^k(\cdot)$ described below.

- All past bids, payments, membership choices, and matches are public.

Each agent’s membership status determines the importance the platform assigns to the agent’s bids, relative to those of others. Suppose, in period $t \geq 1$, agent i from side A submits a positive bid for agent j from side B , whereas agent j submits a negative bid. For given bids (b_{ijt}^A, b_{ijt}^B) , a higher status of agent i implies a higher score for the match (i, j) , thus tilting the allocation in favor of agent i . Symmetrically, a higher status for agent j reduces the score for the match (i, j) , thus tilting the allocation in favor of agent j . A higher status thus grants an agent preferential treatment in all subsequent auctions, both with respect to the competition the agent faces from other agents from his own side (when the capacity constraints are binding) and with respect to the competition the agent faces from agents from the opposite side, shifting the matching to his benefit.

We now turn to the payment rule. Let χ be the matching rule selecting in each period the matches maximizing the sum of the bilateral scores, as described above.¹⁹ For any $t \geq 1$, any (θ, b_t) , any weights $\beta \equiv (\beta_l^k)_{l \in N^k}^{k=A, B}$, let

$$w_t \equiv \sum_{i \in N^A} \sum_{j \in N^B} S_{ijt}^\beta \cdot \chi_{ijt} \quad (4)$$

¹⁸The specific tie-breaking rule plays no role in the analysis. For concreteness, assume ties are broken at random, from a uniform distribution, independently over time.

¹⁹To ease the notation, we drop the dependence of χ on the weights β .

denote the period- t flow weighted surplus implemented under the rule χ . Similarly, let $w_t^{-i,k}$ denote the period- t flow weighted surplus, as defined in (4), in a fictitious environment in which agent i from side k is absent (equivalently, in a market in which all bilateral scores involving agent i from side k are identically equal to zero). In each period $t \geq 1$, the payment asked to each agent $i \in N^A$ (an analogous construction holds for each agent j from side B) is given by

$$\psi_{it}^A = \sum_{j \in N^B} b_{ijt}^A \cdot \chi_{ijt} - \frac{w_t - w_t^{-i,A}}{\beta_i^A(\theta_i^A)}. \quad (5)$$

In words, agent i is asked to make a payment that equals the total flow value the agent derives from all the matches implemented in period t in which the agent is involved (this value is computed as if the agent's own bids reflected the agent's true match values), net of a discount that is proportional to the agent's flow marginal contribution to total weighted surplus, with a coefficient of proportionality inversely determined by the agent's membership status.

As in the VCG auctions, the period- t payments, $t \geq 1$, reflect the flow externalities the agents impose on other agents (from both sides of the market), as well as the costs they impose on the platform. Such externalities may be positive or negative, and therefore the payments may be positive or negative. For example, if an agent is valued highly by the agents he is matched to, he may receive a positive transfer from the platform, reflecting cross-subsidization. Contrary to standard VCG auctions, however, such externalities are calculated with respect to the "weighted surplus" associated with each of the matches, with the weights β reflecting the importance the platform assigns to the agents' bids, which in turn may depend on the platform's objective. Furthermore, multiple agents (from both sides of the market) may be charged for the same externality they impose on others. The following example illustrates:

Example 1 (Many-to-Many Matching with Aggregate Capacity Constraints). Suppose the only relevant capacity constraint is the aggregate one, i.e., $M < n^A \cdot n^B$, but $m_i^k \geq n^{-k}$, all $i = 1, \dots, n^k$, $k = A, B$. Take any agent i from side A (a similar description applies to the side- B agents). In any period in which agent i is matched to some agent from side B his total period- t payment to the platform is equal to

$$\psi_{it}^A = \frac{1}{\beta_i^A(\theta_i^A)} \left(\mathbf{B}_{it}^A(K) + \sum_{j \in N^B} (c_{ijt} - \beta_j^B(\theta_j^B)b_{ijt}^B)\chi_{ijt} \right),$$

where $K \leq n^B$ is the number of agents from side B agent i is matched to, and $\mathbf{B}_{it}^A(K)$ is the sum of the K highest nonnegative scores among the pairs that are unmatched in period t and

that do not include agent i . In the special case in which the platform's costs are identically equal to zero for all matches and $\beta_j^k(\theta_j^k) = 1$, all $\theta_j^k \in \Theta_j^k$, $j = 1, \dots, n^k$, $k = A, B$, the above payments reduces to

$$\psi_{it}^A = \left(\sum_{l \in N^A \setminus \{i\}} \sum_{j \in N^B} (b_{ljt}^A + b_{ljt}^B) \chi_{ljt}^{-i,A} - \sum_{j \in N^B} b_{ijt}^B \chi_{ljt} \right),$$

where

$$\chi_{ljt}^{-i,A} \in \arg \max_{x \in [0,1]^{N^A \times N^B}} \left\{ \sum_{l \in N^A} \sum_{j \in N^B} (b_{ljt}^A + b_{ljt}^B) x_{ljt} : \sum_{l \in N^A} \sum_{j \in N^B} x_{ljt} \leq M, x_{ij} = 0 \text{ all } j \in N^B \right\}$$

is the rule selecting the matches that maximize the some of the bilateral scores in the absence of agent i .

The description of the payment rule is completed by the specification of the tariff the platform uses in period 0 to charge the agents for the selected membership status. The latter is given by

$$\psi_{i0}^A = \theta_i^A D_i^A(\theta) - \int_{\theta_i^A}^{\theta_i^A} D_i^A(y, \theta^{-i,A}) dy - \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \psi_{it}^A \mid \theta \right] - L_i^A, \quad (6)$$

where $\theta^{-i,A}$ is the profile of membership statuses for all agents excluding agent i from side A ,²⁰

$$D_i^A(\theta) \equiv \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j \in N^B} \varepsilon_{ijt}^A \chi_{ijt} \mid \theta \right] \quad (7)$$

is the “quality” of the matches agent i expects from joining the platform, when the profile of membership statuses is θ and all agents are expected to bid their true match values v_{ljt}^k in all periods, and L_i^A is a scalar. The term $\int_{\theta_i^A}^{\theta_i^A} D_i^A(y, \theta^{-i,A}) dy + L_i^A$ in ψ_{i0}^A is a discount whose role is to guarantee that the agent has incentives to join in period zero and select the status designed for his true vertical type. Note that, according to (6), the fee the platform charges to the agent in period 0 depends on the entire profile of membership statuses selected by all agents. In other words, the price for status depends on the aggregate demand of status from each side of the market. Such dependence plays a role analogous to that of “insulating tariffs”

²⁰The same notation will be used also for any other vector. Given any $z = (z_i^k)_{i=1, \dots, n^k}^{k=A,B}$, $z^{-i,k}$ will denote the vector obtained from z by deleting z_i^k . Similarly, given $a_t = (a_{ijt}^k)_{i,j \in N^A \times N^B}^{k=A,B}$, $a_t^{-i,k}$ will denote the vector obtained from a by eliminating $(a_{ijt}^k)_{j=1, \dots, n-k}$.

in two-sided markets (see, e.g., Weyl, 2010): it guarantees that all agents find it optimal to join the platform in period zero, irrespective of their beliefs about the types of other agents. Such insulating tariffs can always be replaced by (perhaps more familiar) tariffs in which the price each agent pays in period 0 depends only on the agent's own membership status, by taking expectations over the types of other agents.

4 Equilibrium Bidding

Definition 2 (Truthful Strategies). A strategy profile $\sigma = (\sigma_l^k)_{l \in N^k}^{k=A,B}$ for the above matching auctions is *truthful* if at $t = 0$ each agent selects the membership status designed for his true vertical type and, at each period $t \geq 1$, for any history, each agent's bids coincide with the agent's true match values, i.e., $b_{ijt}^k = v_{ijt}^k = \theta_l^k \cdot \varepsilon_{ijt}^k$, all $(i, j) \in N^A \times N^B$, $k = A, B$. A truthful equilibrium is an equilibrium in which the strategy profile is truthful.

We have the following result:

Theorem 1. *Any matching auction in which the scores are given by (3), with arbitrary weights β , and the payments are given by (5) and (6), with L_l^k large enough, all $l \in N^k$, $k = A, B$, admits an equilibrium in which all agents participate in each period and follow truthful strategies. Furthermore, such truthful equilibria are periodic ex-post (that is, the agents' strategies are sequentially rational, regardless of the agents' beliefs about other agents' past and current types).*

Heuristic Proof. The formal proof is in the Appendix. Here we illustrate the key ideas in an heuristic way. We organize the arguments in two steps. The first step establishes that, when the payments are the ones in (5) above, in the continuation game that starts with period $t \geq 1$, irrespective of past behavior, all agents find it optimal to participate in the auctions and bid truthfully. Step 2 shows that, when the membership fees are the ones in (6), all agents find it optimal to join the platform in period zero and select the membership status designed for their true vertical type, again regardless of their beliefs about other agents' types.

Step 1. The payments in (5) are designed so that, irrespective of past histories, at each period $t \geq 1$, an agent who expects all other agents to participate and bid truthfully from period t onwards has incentives to do the same. To see this, observe that, under such payments, given the selection of the membership statuses at $t = 0$, incentives separate over time, meaning that the period- t bids that maximize each agent's continuation payoff (net of the payments) are invariant to the bids the same agent submitted in previous periods and in the bids he will submit in subsequent periods. Furthermore, under such payments, the agent's period- t flow payoff (once again, net of the flow payments) is proportional to his flow marginal

contribution to *weighted surplus* (with the latter defined as in (4), and with the coefficient of proportionality given by $1/\beta_t^k(\theta_t^k)$). Because, given the received bids, the matches implemented by the auction maximize the flow total weighted surplus after all histories (including those off the equilibrium path), agents have incentives to remain in the auctions and bid truthfully after all histories, irrespective of the beliefs they may have about past and current types of other agents. In other words, participating and following truthful strategies constitutes a periodic ex-post continuation equilibrium, after any period- t history, $t \geq 1$. This step is standard and follows from essentially the same arguments establishing the optimality of truthful bidding in standard VCG mechanisms, with the only difference being that the proposed matches maximize weighted surplus instead of true surplus.

Step 2. Showing that all agents find it optimal to join the platform in period zero and select the membership status designed for their true vertical type is more delicate. The proof in the Appendix shows that, when the membership fees are the ones in (6), and all agents follow truthful strategies, each agent i from side A whose true vertical type is θ_i^A and who believes his opponents' vertical types to be $\theta^{-i,A}$ expects a total payoff equal to $\int_{\theta_i^A}^{\theta_i^A} D_i^A(y, \theta^{-i,A}) dy + L_i^A$, where $D_i^A(\theta_i^A, \theta^{-i,A})$ is the average match quality expected by the agent (as defined in (7)) when the profile of true vertical types is $(\theta_i^A, \theta^{-i,A})$ and all agents, including i , follow truthful strategies (clearly, the same representation applies to any side- B agent). Provided that L_i^A is sufficiently large to compensate for the possibly negative match quality expected by some types below θ_i^A , the agent then has incentives to participate at period zero. Establishing that the agent finds it optimal to select the membership status designed for his true vertical type is more demanding and is done by showing that the match quality expected by an agent of true vertical type θ_i^A who selects the membership status for type $\hat{\theta}_i^A$ is nondecreasing in $\hat{\theta}_i^A$, irrespective of the profile θ of true vertical types. In the Appendix, we also show that, under truthful strategies, agents with a higher vertical type who select the membership status designed for their true type expect higher match quality than lower types. This second monotonicity is not crucial for incentive compatibility but guarantees that, when lower types find it optimal to participate, so do higher types. Such property plays an important role in the design of profit-maximizing auctions, as we show in the next section. ■

Note that, to show that all agents find it optimal to join the platform in period zero and select the membership status designed for their true vertical type (regardless of their beliefs about other agents' types), one cannot simply use standard VCG arguments, for the selection of the membership status at $t = 0$ is precisely what endogenizes the weights used in the subsequent auctions. Incentive compatibility at $t = 0$ is thus established using arguments similar to those in Pavan et al. (2014), but adapted to account for the fact that the private information

the agents expect to receive in each period following the initial one is multidimensional.

5 Profit and Welfare Maximization

We now show that the matching auctions introduced in Section 3 include a subclass that maximizes the platform's profits over all possible mechanisms as well as a class that maximizes welfare over all possible mechanisms.²¹ Equipped with these results, we then examine the distortions in the provision of matching services due to profit maximization.

In this section, we make the additional assumptions that, for all $l \in N^k$, $k = A, B$, the function $[1 - F_l^k(\theta_l^k)] / f_l^k(\theta_l^k)\theta_l^k$ is non-increasing with $f_l^k(\theta_l^k)\theta_l^k > 1$.²² In the next two theorems, we highlight the role of the weights β by denoting by $D_l^k(\theta; \beta)$ the intertemporal match quality expected by agent l from side k when the scores in (3) use the weights β . We then have the following result:

Theorem 2 (Profit-Maximizing Auctions). *Let β^P be the weights given by²³*

$$\beta_l^{k,P}(\theta_l^k) \equiv 1 - \frac{1 - F_l^k(\theta_l^k)}{f_l^k(\theta_l^k)\theta_l^k}, \text{ all } l \in N^k, k = A, B. \quad (8)$$

Suppose that $D_l^k(\theta_l^k, \theta^{-l,k}; \beta^P) \geq 0$, all $l \in N^k$, $k = A, B$, $\theta^{-l,k} \in \Theta^{-l,k}$. The matching auction in which (a) the scores are given by (3) with weights β^P , and (b) the payments are given by (5) and (6) with weights β^P and with $L_l^k = 0$, all $l \in N^k$, $k = A, B$, are profit maximizing.

The proof in the Appendix is in three steps. First, we show that, given any matching mechanism $\Gamma \equiv (\mathcal{M}, \mathcal{S}, \chi, \psi, \rho)$ and any *Bayes Nash equilibrium* (and hence any PBE) σ of the game induced by Γ , the period-0 interim expected payoff of each agent $l \in N^k$ from each

²¹A matching mechanism will be referred to as *profit maximizing* if it is feasible and admits a PBE under which the platform's profits are as high as under any other PBE of any other feasible mechanism. A *welfare-maximizing* mechanism is defined in a similar way, but with welfare replacing profits in the platform's objective.

²²The first part of the assumption guarantees that the agents' "virtual" vertical types $\theta_l^k - [1 - F_l^k(\theta_l^k)] / f_l^k(\theta_l^k)$ are non-decreasing in the true types and that the weights β^P in the next theorem are also non-decreasing. The second part is added to guarantee that the virtual vertical types are strictly positive. Given the multiplicative structure of the match values in (1), this assumption guarantees that the virtual match values $[\theta_l^k - [1 - F_l^k(\theta_l^k)] / f_l^k(\theta_l^k)] \varepsilon_{ijt}^k$ respect the same rankings as the true ones, $v_{ijt}^k = \theta_l^k \varepsilon_{ijt}^k$, thus avoiding confusion in the interpretation.

²³Note that, under the maintained assumptions on the F distributions, these weights are non-decreasing, strictly positive, and bounded, as required by the definition of the matching auctions in Section 3.

side $k = A, B$ is given by²⁴

$$U_l^k(\theta_l^k) = U_l^k(\underline{\theta}_l^k) + \int_{\underline{\theta}_l^k}^{\theta_l^k} \mathbb{E} [D_l^k(\theta; \hat{\chi})|y] dy, \quad (9)$$

where $D_l^k(\theta; \hat{\chi})$ is the expected match quality when, in each period, given the state (θ, ε^t) , the selected matches are the ones under the matching rule $\hat{\chi}$ obtained by compounding χ with σ , i.e., $\hat{\chi}_t(\theta, \varepsilon^t) = \chi_t(\sigma(\theta, \varepsilon^t))$, and where the expectation in (9) is with respect to the entire profile of vertical types $\theta = (\theta_l^k, \theta^{-l,k})$, given the agent's own vertical type, θ_l^k .

Next, we use the above representation of the agents' equilibrium interim expected payoffs to show that, given any mechanism Γ and any BNE σ of Γ , the platform's profits are given by the following weighted surplus function

$$\mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{i \in N^A} \sum_{j \in N^B} \left(\beta_i^{A,P}(\theta_i^A) \theta_i^A \varepsilon_{ijt}^A + \beta_j^{B,P}(\theta_j^B) \theta_j^B \varepsilon_{ijt}^B - c_{ijt} \right) \hat{\chi}_{ijt}(\theta, \varepsilon^t) \right] - \sum_{k=A,B} \sum_{l \in N^k} U_l^k(\underline{\theta}_l^k), \quad (10)$$

for which the first term is only a function of the matching rule, and with the weights β^P as given in (8).

The optimality of the matching auctions of Theorem 2 then follows from the fact, under the truthful equilibria of the proposed auctions, (a) the induced state-contingent matches maximize the first component of the weighted surplus function (10), and (b) the period-0 participation constraints of the lowest vertical types are binding. That the lowest vertical type of each agent expects a nonnegative match quality in the matching auctions with weights β^P (i.e., that $D_l^k(\underline{\theta}_l^k, \theta^{-l,k}; \beta^P) \geq 0$), together with the fact that, under truthful strategies, match quality $D_l^k(\cdot, \theta^{-l,k}; \beta^P)$ is non-decreasing in each agent's vertical type θ_l^k (this is one of the two key monotonicity properties discussed after Theorem 1), in turn guarantees that participation in period zero is optimal for all agents. In fact, in the proof of Theorem 2, we show that payoffs satisfy a stronger envelope condition, which guarantees the optimality of period-zero participation, irrespective of the agents' beliefs about other agents' types.

The conditions in Theorem 2 can be relaxed by requiring that the match quality $\mathbb{E} [D_l^k(\theta; \beta^P) | \underline{\theta}_l^k]$ expected, on average, by the lowest vertical type of each agent is non-negative, where the average is over other agents' vertical types $\theta^{-l,k}$. These conditions are vacuously satisfied if horizontal types are nonnegative, that is, if no agent dislikes interacting with any other agent

²⁴By period-0 interim expected payoff, we mean the payoff the agent expects, under the equilibrium strategy profile σ , in the game induced by Γ , when his period-0 vertical type is equal to θ_l^k .

from the opposite side. Interestingly, the truthful equilibria of the matching auctions of Theorem 2 maximize the platforms' profits over all BNE (not just PBE) of all feasible mechanisms.

We now turn to welfare maximization.

Theorem 3 (Welfare-Maximizing Auctions). *Let β^W be the weights given by $\beta_l^{k,W}(\theta_l^k) = 1$, all θ_l^k , $l \in N^k$, $k = A, B$.*

(i) *The matching auctions in which (a) the scores are given by (3), with weights β^W , and (b) the payments are given by (5) and (6), with weights β^W and with L_l^k large enough, all $l \in N^k$, $k = A, B$, are welfare maximizing.*

(ii) *Suppose $D_l^k(\underline{\theta}_l^k, \theta^{-l,k}; \beta^W) \geq 0$, all $l \in N^k$, $k = A, B$, and $\theta^{-l,k} \in \Theta^{-l,k}$. The matching auctions with weights β^W and payments given by (5) and (6), with $L_l^k = 0$, all $l \in N^k$, $k = A, B$, admit periodic ex-post equilibria in which agents participate and follow truthful strategies at all histories. Furthermore, the truthful equilibria of these auctions maximize the platform's profits over all BNE of all feasible mechanisms implementing welfare-maximizing matches and inducing the agents to join the platform in period zero.*

The result in part (i) follows directly from the fact that, when the weights are given by β^W , the matches sustained under the truthful strategies of the proposed matching auctions maximize welfare after each history. The conditions in part (ii) in turn guarantee that, when the payments are as in (5) and (6) with weights β^W and with $L_l^k = 0$, all $l \in N^k$, $k = A, B$, each agent finds it optimal to participate in period zero, irrespective of his beliefs about other agents' types. This follows from the fact that each agent's expected payoff at the time of joining is given by (9) along with the assumption that the match quality $D_l^k(\underline{\theta}_l^k, \theta^{-l,k}; \beta^W)$ expected by the lowest vertical type of each agent is non-negative, for all possible $\theta^{-l,k}$, and the fact that, given $\theta^{-l,k}$, $D_l^k(\cdot, \theta^{-l,k}; \beta^W)$ is non-decreasing in $\theta^{-l,k}$, as established in Lemma 2 in the Appendix.

That the truthful equilibria of the proposed auctions maximize the platform's profits over all BNE of all feasible mechanisms implementing welfare-maximizing matches and inducing the agents to participate in period zero follows from arguments similar to those establishing the optimality of the matching auctions in Theorem 2 above. To see this, note that the platform's profits under any BNE (and hence under any PBE) of any feasible mechanism implementing the welfare-maximizing matches and inducing the agents to participate at $t = 0$ are given by (10), with $\hat{\chi}$ representing the efficient assignment rule, and with $U_l^k(\theta_l^k) \geq 0$, all $l \in N^k$, $k = A, B$. Under the truthful equilibria of the proposed auctions, the interim payoff $U_l^k(\underline{\theta}_l^k)$ expected by the lowest vertical type of each agent is exactly equal to zero. Hence, the platform's profits under the truthful equilibria of the proposed matching auctions are at least

as high as under any BNE of any feasible mechanism inducing the agents to participate in period zero and implementing the efficient matches.

We conclude this section by discussing the distortions brought in by market power under profit maximization. Let $\chi^P \equiv (\chi_t^P(\theta, \varepsilon))_{t=1}^\infty$ and $\chi^W \equiv (\chi_t^W(\theta, \varepsilon))_{t=1}^\infty$ denote the profit- and welfare-maximizing matching allocations, under the truthful equilibria of the auctions of Theorems 2 and 3, respectively. Because (θ, ε) are exogenous, we then drop them from the arguments of χ^P and χ^W below to ease the notation. We then have the following result:²⁵

Theorem 4 (Distortions). *Suppose that all agents derive a nonnegative utility from interacting with all other agents from the opposite side (formally, $\varepsilon_{ijt}^k \geq 0$, all $(i, j) \in N^A \times N^B$, $k = A, B$, $t \geq 1$).*

(1) *If none of the capacity constraints binds (i.e., $M \geq n^A \cdot n^B$, and $m_l^m \geq n^{-k}$, all $l \in N^k$, $k = A, B$), then, for all $(i, j) \in N^A \times N^B$, all $t \geq 1$, $\chi_{ijt}^P = 1 \Rightarrow \chi_{ijt}^W = 1$.*

(2) *If only the platform's aggregate capacity constraint is potentially binding (i.e., $M < n^A \cdot n^B$, but $m_l^k \geq n^{-k}$, all $l \in N^k$, $k = A, B$), then, for all $t \geq 1$, $\sum_{(i,j) \in N^A \times N^B} \chi_{ijt}^W \geq \sum_{(i,j) \in N^A \times N^B} \chi_{ijt}^P$. Nonetheless, certain matches active under profit-maximization need not be active under welfare maximization.*

(3) *If some of the individual capacity constraints are potentially binding (i.e., if $m_l^m < n^{-k}$, for some $l \in N^k$, $k = A, B$), then, for all $t \geq 1$, $\sum_{(i,j) \in N^A \times N^B} \chi_{ijt}^P > 0$ implies that $\sum_{(i,j) \in N^A \times N^B} \chi_{ijt}^W > 0$.*

As in other screening problems, distortions are introduced under profit maximization to reduce the agents' information rents (that is, the surplus the platform must leave to the agents to induce them to reveal their private information). When all agents value positively interacting with all other agents from the opposite side and none of the capacity constraints binds, in each period, a profit-maximizing platform induces fewer interactions than a welfare-maximizing one. In particular, any match that is active under profit-maximization is also active under welfare maximization. This is because, when none of the capacity constraints binds, both under profit and under welfare maximization, the platform implements in each period all matches for which the score is non-negative. The result then follows from the fact that, at each period, and for each match, the score under profit maximization is smaller than under welfare maximization due to the handicaps $[1 - F_l^k(\theta_l^k)]/f_l^k(\theta_l^k)\theta_l^k$, $k = A, B$, the platform applies to the scores to account for the cost of leaving rents to agents with a vertical type above θ_l^k .

²⁵To facilitate the comparison between profit and welfare maximization, we assume that, in each period, both the profit- and the welfare-maximizing auctions match pairs for which the bilateral score is zero when the capacity constraints are not binding.

When, instead, some of the capacity constraints are potentially binding, certain matches active under profit maximization need not be active under welfare maximization. What remains true though is that, when the only potentially binding constraint is the platform's, at any point in time, the total number of matches under welfare maximization is at least as high as under profit maximization. This property follows again from the fact that each score under welfare maximization is at least as large as under profit maximization, along with the fact that the total number of matches active under both profit and welfare maximization is the minimum between the number of matches for which the score is non-negative and the platform's aggregate capacity M . That some of the matches active under profit maximization need not be active under welfare maximization follows from the fact the ranking of the scores when the weights are given by β^P need not coincide with the ranking of the scores when the weights are given by β^W .

Interestingly, the property that, at any given point in time, the total number of matches active under welfare maximization is at least as large as under profit maximization need not extend to markets in which some of the individual capacity constraints potentially bind. To see this, suppose there are 2 agents on each side of the market, matching is one-to-one (i.e., $m_i^k = 1$, $i = 1, 2$, $k = A, B$) and $M > 4$, so that the only relevant capacity constraints are the individual ones. Let S_{ijt}^P and S_{ijt}^W be the period- t scores, under the weights β^P and β^W , respectively. Further suppose that $S_{22t}^W < 0$, and that $S_{ijt}^P > 0$ for any $(i, j) \neq (2, 2)$. Lastly assume that $S_{11t}^W > S_{12t}^W + S_{21t}^W$, whereas $S_{11t}^P < S_{12t}^P + S_{21t}^P$. Then, despite all scores being higher under welfare maximization than under profit maximization, the welfare-maximizing auction implements a single match in period t , the one corresponding to $(1, 1)$, whereas the profit-maximizing auction implements two matches in the same period, $(1, 2)$ and $(2, 1)$. The reason is that, once the match $(1, 1)$ is formed, the matches $(1, 2)$ and $(2, 1)$ become unfeasible, due to the individual capacity constraints. What remains true though is that, at any period in which matching is not completely shut down under profit maximization (i.e., at least one match is active), matching is also not completely shut down under welfare maximization (part (iii) in the Theorem).

The above conclusions need not extend to settings in which agents dislike certain interactions (formally, when horizontal types may take on negative values for certain pairs). In such settings, a profit-maximizing platform may induce an inefficiently *high* number of matches within each period, and, over time, each pair may experience a larger number of matches under profit maximization than under welfare maximization, *irrespective of the capacity constraints*. Formally, when certain interactions may generate negative match values, a pair's score under profit maximization may be greater than its counterpart under welfare maximization, which

implies that the total number of interactions under profit maximization may exceed the efficient level. The reason why a profit-maximizing platform may induce an inefficiently high number of interactions is that this may discourage the agents from purchasing a lower membership status. By locking those agents selecting a low status (equivalently, claiming to have a low vertical type) into unpleasant interactions, the platform makes it costly for those agents with a high vertical type to pretend to have a low type. In turn, this permits the platform to extract more surplus from those high-type agents. The following example illustrates.

Example 2 (Upward Distortions under Negative Values). Consider the following environment where $N^A = N^B = \{1\}$, $M = 1$, and $c_{11t} = 0$, all $t \geq 1$. The vertical types are given by $\Theta_1^B = \{1\}$ and $\Theta_1^A = [1 + \varsigma, 2 + \varsigma]$, $\varsigma > 0$, with F_1^A uniform over Θ_1^A . At each period $t \geq 1$, regardless of past realizations, $\varepsilon_{11t}^B = 1$, whereas ε_{11t}^A is drawn uniformly from $\{-3, +3\}$. Suppose the realized vertical type of agent 1 from side A is equal to $1 + \varsigma$, in which case the weights used under profit maximization to scale the two agents' bids are given by $\beta_1^{A,P}(\theta_1^A) = \varsigma$ and $\beta_1^{B,P}(\theta_1^B) = 1$. Furthermore, consider a realized sequence $(\varepsilon_{11t}^A)_{t=1}^\infty$ of horizontal types for agent 1 from side A such that $\varepsilon_{11t}^A = -3$, all $t \geq 1$. Then, for sufficiently small ς and δ , the pair is matched in each period under profit maximization, despite matching being inefficient.

Finally, note that the familiar result of “no distortion at the top” from standard screening problems does not apply to a matching environment. A profit-maximizing platform may distort the matches of all agents, including those “at the top” of the distribution, for whom the vertical type is the highest. The reason is that, contrary to standard screening problems in which the cost of procuring inputs is exogenous, in a matching market, the cost of “procuring” agents-inputs from the opposite side of the market is endogenous and is higher than under welfare maximization, due to the informational rents the platform must provide to such agents-inputs to induce them to reveal their private information.

The above results bear certain implications for government intervention in matching markets, a topic that is receiving a lot of attention in recent years. In markets in which capacity constraints are unlikely to be binding and agents are unlikely to suffer losses from interacting with other agents, platforms should be encouraged to implement more matches. Importantly, even if platforms could be induced to run welfare-maximizing auctions (more generally, to implement the welfare-maximizing matches), such auctions are not guaranteed to yield positive profits to the platforms (despite the fact that they minimize the platforms' losses over all mechanisms implementing the welfare-maximizing matches, as established in part (ii) of Theorem 3). As a result, the government may need to subsidize such markets. On the other hand, in markets where cross-subsidization and/or binding capacity constraints are prominent features, it may be desirable to discourage certain interactions.

6 Endogenous Processes

We now show how matching auctions similar to those introduced above but with “forward-looking” scores can be used in certain markets in which the evolution of the agents’ match values is endogenous. To this purpose, consider a setting identical to the one in the previous sections, but where the kernels governing the evolution of the match values now depend on past allocations and satisfy the following properties (as in the baseline model, we focus on a representative agent i from side A with the understanding that the same properties apply also to each agent l from side B): (1) for any $(j, j') \in N^B$, $j' \neq j$, given the matches x , the sequence of horizontal types $\varepsilon_{ij}^A \equiv (\varepsilon_{ij^t}^A)_{t=1, \dots, \infty}$ is drawn independently from the sequence $\varepsilon_{ij'}^A \equiv (\varepsilon_{ij'^t}^A)_{t=1, \dots, \infty}$; (2) whenever $x_{ij^{t-1}} = 1$, the dependence of the kernel $G_{ij^t}^A(\varepsilon_{ij^t}^A \mid \varepsilon_{ij^{t-1}}^A, x_{ij}^{t-1})$ on x_{ij}^{t-1} is only through $\sum_{s=1}^{t-1} x_{ijs}$; (3) whenever, instead, $x_{ij^{t-1}} = 0$, $G_{ij^t}^A$ is a Dirac measure at $\varepsilon_{ij^t}^A = \varepsilon_{ij^{t-1}}^A$, i.e., $G_{ij^t}^A(\varepsilon_{ij^t}^A \mid \varepsilon_{ij^{t-1}}^A, x_{ij}^{t-1}) = \mathbf{1}_{\{\varepsilon_{ij^t}^A \geq \varepsilon_{ij^{t-1}}^A\}}$; (4) there exists a sequence $(\omega_{ijs}^A)_{s=1, \dots, \infty} \in \mathbb{R}^\infty$ drawn from an exogenous distribution, such that, for any number R_{ij} of past interactions between agent $i \in N^A$ and agent $j \in N^B$, $\varepsilon_{ij^t}^A$ is given by a deterministic function of $(\omega_{ijs}^A)_{s=1}^{R_{ij}}$, uniformly over t .²⁶

The above assumptions imply that match values are drawn independently across relationships and change only upon interacting with partners. Furthermore, the processes governing the agents’ match values are Markov time-homogeneous and their dependence on past matches is only through the number of past interactions. As we show below, when paired with appropriate conditions on the capacity constraints, such assumptions guarantee that the match dynamics under both profit and welfare maximization continue to be implementable by matching auctions similar to those in the previous sections but with scores taking the form of indexes (more below).

In addition to the above changes, assume the period- t costs $c_{ij^t}(x_{ij}^{t-1})$ the platform incurs for matching the pair (i, j) also depend on the history of past matches x_{ij}^{t-1} through the number of past interactions $\sum_{s=1}^{t-1} x_{ijs}$.²⁷

The following scenario is consistent with the aforementioned assumptions.

Example 3 (Gaussian Learning). Each agent $i \in N^A$ derives a constant utility $v_{ij^t}^A = \theta_{ij}^A u_{ij}^A$ for interacting with each agent $j \in N^B$. Such utility is unknown to the platform and to all

²⁶In addition to assumptions (1)-(4), as in the baseline model, we continue to assume that, for any $i \in N^A$, there exists a constant $E_i^A > 0$ such that, for any sequence of matches x , $\mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j \in N^B} |\varepsilon_{ij^t}^A| \cdot x_{ij^t} \right] \leq E_i^A$, where the expectation is taken with respect the distribution over \mathcal{E} generated by the kernels G , under the matches x .

²⁷A special case of interest is when the costs vanish after the first interaction, which corresponds to a situation where the only relevant costs are those that the platform incurs to “link” the agents.

agents. Agent i starts with a prior belief that $u_{ij}^A \sim N(\varepsilon_{ij1}^A, \tau_{ij}^A)$, where the variance τ_{ij}^A is common knowledge but where the initial prior mean ε_{ij1}^A is the agent's private information. The agent's prior mean ε_{ij1}^A is drawn from a distribution G_{ij1}^A . Each time agent i is matched to agent j , agent i receives a conditionally i.i.d. private signal $\xi_{ij}^A \sim N(u_{ij}^A, \vartheta_{ij}^A)$ about his idiosyncratic appreciation for agent j , u_{ij}^A , and updates his expectation of u_{ij}^A using Bayes rule.²⁸ In this setting, the vector of period- t horizontal types $\varepsilon_{it}^A \equiv (\varepsilon_{ijt}^A)_{j \in N-k}$ thus corresponds to the collection of the agent's posterior means about the attractiveness of potential partners from the opposite side.

Another scenario consistent with the aforementioned assumptions is the following one:

Example 4 (Preference for Variety or Habit Formation). The value each agent $i \in N^A$ derives from interacting with each agent $j \in N^B$ decreases (alternatively, increases) with the number of past interactions with agent j . Precisely, for all $t \geq 1$, all $(\varepsilon_{ij,t-1}^A, \varepsilon_{ij,t}^A)$, $G_{ij,t}^k(\varepsilon_{ij,t}^A | \varepsilon_{ij,t-1}^A, x_{ij}^{t-1})$ is non-decreasing (alternatively, non-increasing) in $\sum_{s=1}^{t-1} x_{ij,s}$. For example, in the case of preference for variety, agents gradually lose interest in partners they already interacted with, whereas in the case of habit formation, the value each agent assigns to each partner increases with the number of past interactions.²⁹

Now consider matching auctions similar to the ones described in the previous sections but where, in each period $t \geq 1$, each agent $l \in N^k$, from each side $k = A, B$, in addition to submitting bids for each possible partner from the opposite side, is offered the possibility to revise his membership status by selecting $\theta_{lt}^k \in \Theta_l^k$. The reason why the platform allows the agents to revise their status over time, despite the fact that, under the specification in (1), the vertical types are perfectly persistent, is that this favors equilibria in which the agents bid myopically over time, irrespective of past bids and membership selections. In fact, while the agents' bids convey information about the agents' total match values $v_{ij,t}^k$, with endogenous processes, they are not sufficient statistics with respect to the agents' overall private information, when it comes to predicting future values. For instance, a high bid $b_{ij,t}^k$ by agent i for agent j may either reflect a persistent high value for interacting with all agents from the opposite side (i.e., a high vertical type θ_i^k), or a high temporary appreciation for interacting with agent j (i.e., a high horizontal type, $\varepsilon_{ij,t}^k$). When the evolution of the agents'

²⁸Specifically, each signal ξ_{ij}^A can be written as $\xi_{ij}^A = u_{ij}^A + \zeta_{ij}^A$, with the innovations ζ_{ij}^A drawn from a Normal distribution with mean 0 and variance ϑ_{ij}^A , independently from all other random variables.

²⁹A special form of the preference for variety is when each agent demands a fixed number of interactions with each partner, which amounts to assuming that, for each agent $i \in N^A$, each partner $j \in N^B$, there exists a number $\alpha_{ij}^A \in \mathbb{N}$ such that, whenever $\sum_{s=1}^{t-1} x_{ij,s} > \alpha_{ij}^A$, $G_{ij,t}^A(\varepsilon_{ij,t}^A | \varepsilon_{ij,t-1}^A, x_{ij}^{t-1}) = 1$ for all $\varepsilon_{ij,t}^A \geq 0$. Note that a market where each agent is able to provide only up to a fixed number of services to each partner is also a special case of the above specification.

match values is endogenous, being able to collect all the information that is necessary to predict the evolution of the agents' match values is key to a proper selection of the matches under profit- and welfare-maximizing auctions, as we show next.

Let $\lambda_{ij}|\theta_t, b_t, x^{t-1}$ denote the stochastic process over pair (i, j) 's current and future match values (v_{ijs}^A, v_{ijs}^B) that one obtains under any matching rule that matches the pair (i, j) at all periods $s \geq t$, when the pair's vertical types (θ_i^A, θ_j^B) are the ones revealed by the period- t membership choices θ_t , and when the horizontal types are given by

$$\varepsilon_{ijt}^k = \begin{cases} b_{ijt}^k/\theta_{it}^k & \text{if } b_{ijt}^k/\theta_{it}^k \in \mathcal{E}_{ijt}^k \\ \arg \min_{\hat{\varepsilon}_{ijt}^k \in \mathcal{E}_{ijt}^k} \{|b_{ijt}^k/\theta_{it}^k - \hat{\varepsilon}_{ijt}^k|\} & \text{otherwise,} \end{cases} \quad (11)$$

for $k = A, B$.

Definition 3 (Index Scores). An index scoring rule $S^{I;\beta}$ (with weights β) is one in which, for each $t \geq 1$, each pair $(i, j) \in N^A \times N^B$, each (θ_t, b_t, x^{t-1}) , the period- t score for the (i, j) -match is given by

$$S_{ijt}^{I;\beta} \equiv \max_{\tau} \left\{ \frac{\mathbb{E}^{\lambda_{ij}|\theta_t, b_t, x^{t-1}} \left[\sum_{s=t}^{\tau} \delta^{s-t} (\beta_i^A(\theta_{i0}^A)v_{ijs}^A + \beta_j^B(\theta_{j0}^B)v_{ijs}^B - c_{ijs}) \right]}{\mathbb{E}^{\lambda_{ij}|\theta_t, b_t, x^{t-1}} \left[\sum_{s=t}^{\tau} \delta^{s-t} \right]} \right\}, \quad (12)$$

where τ denotes a stopping time.

The period- t index score $S_{ijt}^{I;\beta}$ for the match between agent i from side A and agent j from side B thus corresponds to a Gittins index for a process for which the ‘‘rewards’’ are given by the flow weighted total surplus of the match (i, j) , with the weights given by the functions β of the agents' period-0 membership statuses. Note that each score $S_{ijt}^{I;\beta}$ depends on x^{t-1} only through the total number of past interactions $\sum_{s=1}^{t-1} x_{ijs}$ between the pair (i, j) . Also note that, contrary to the myopic scores in the auctions in the previous sections, the indexes $S_{ijt}^{I;\beta}$ depend on the entire history of past and current membership choices. However, while the dependence on the period-0 and on the period- t choice is direct, the dependence on other periods' choices is only through the number of past interactions.

Now consider auctions that select in each period the matches for which the sum of the index scores is the highest, subject to individual and aggregate capacity constraints, and where the payments are defined as follows. Let

$$W_t \equiv \mathbb{E}^{\lambda[x^{I;\beta}]|\theta_t, b_t, x^{t-1}} \left[\sum_{s=t}^{\infty} \delta^{s-t} \sum_{i \in N^A} \sum_{j \in N^B} (\beta_i^A(\theta_{i0}^A)v_{ijs}^A + \beta_j^B(\theta_{j0}^B)v_{ijs}^B - c_{ijs}) \cdot \chi_{ijs} \right] \quad (13)$$

denote the continuation weighted surplus, that is, the discounted present value of the weighted sum of match values, net of the platform's costs, when all agents follow truthful strategies, i.e., when they bid truthfully their myopic values v_{ijs}^k and select the membership status corresponding to their true vertical types, at all periods $s \geq t$. Here $\lambda[\chi^{I;\beta}]|\theta_t, b_t, x^{t-1}$ denotes the stochastic process over (θ_s, v_s, x^{s-1}) , $s \geq t$, when the selected period- t membership statuses are θ_t , the period- t bids are b_t , all agents follow truthful strategies from period $s > t$ onwards, the true vertical types are the ones corresponding to the selected period- t membership statuses (i.e., $\theta_l^k = \theta_{lt}^k$, $l \in N^k$, $k = A, B$), and the true period- t horizontal types are given by (11).

Similarly, let $W_t^{-l,k}$ denote the continuation weighted surplus, as defined in (13), in a fictitious market without agent l from side k . Next, let $R_{it}^k \equiv W_t - W_t^{-l,k}$ denote the *contribution* of agent $l \in N^k$ to the continuation weighted surplus and

$$r_{it}^k \equiv R_{it}^k - \delta \mathbb{E}^{\lambda[\chi^{I;\beta}]|\theta_t, b_t, x^{t-1}} [R_{it+1}^k] \quad (14)$$

the corresponding *flow marginal contribution*. In each period $t \geq 1$, the payment asked to each agent $i \in N^k$ is given by

$$\psi_{it}^k = \sum_{j \in N^{-k}} b_{ijt}^k \cdot \chi_{ijt} - \frac{1}{\beta_i^k(\theta_{i0}^k)} r_{it}^k. \quad (15)$$

The above payments are similar to those in Bergemann and Valimaki (2010) but adapted to account for the fact that the platform's objective may differ from profit maximization (see also Kakade et al. (2013) for a similar construction). The period-0 payments, instead, continue to be given by (6), as in the baseline model, but with $D_i^k(\theta; \beta)$ now denoting the expected match quality in the auctions where the scores are the ones in (12) and the expectations are computed under the endogenous distribution $\lambda[\chi^{I;\beta}]|\theta_0$.³⁰

We then have the following result:

Theorem 5 (Endogenous Processes). *Suppose that either $M = 1$, or none of the capacity constraints binds (i.e., $M \geq n^A \cdot n^B$, and $m_l^m \geq n^{-k}$, all $l \in N^k$, $k = A, B$). Then conclusions analogous to those in Theorems 1-3 apply to the matching auctions where the scores are given by (12) and the payments are given by (6) and (15). Furthermore, the same conclusions as in parts (1) and (2) of Theorem 4 hold (with $M = 1$, for part 2).*

The proof follows from steps similar to those establishing Theorems 1-4, but adjusted to

³⁰This is the endogenous distribution over match values that obtains when the scores are those in (12), agents follow truthful strategies from period $t = 1$ onwards, and the true vertical types are those corresponding to the period-0 membership statuses, θ_0 .

account for the fact that the maximization of continuation weighted surplus is not separable over time.

That, under the proposed auctions, at each period after the initial one, agents have incentives to remain in the mechanism, select the membership status designed for their true vertical type, and bid truthfully follows from arguments similar to those in Bergemann and Valimaki (2010) and Kakade et al. (2013). The key difficulty is showing that they also have the right incentives to participate in period zero and select the period-0 membership status designed for their true types. As in the baseline model, this is done by showing that, under truthful bidding, the match quality each agent expects when joining the platform is non-decreasing in both the selected membership status, for given true vertical type, and in the agent's true vertical type when the latter coincides with the selected one.

With endogenous processes, the problem of maximizing continuation weighted surplus is a *multi-armed bandit problem*. Such a problem is known to admit an index solution only under relatively stringent conditions (namely, when rewards evolve independently across arms, arms are frozen when not activated, and a single arm, or any number of arms, can be activated in each period). Despite the obvious limitations imposed by the conditions guaranteeing the optimality of index scores, we believe the matching auctions introduced above capture many of the relevant trade-offs that platforms face in the design of their dynamic matching protocols.

A key difficulty when agents learn endogenously their match values by interacting with other agents is that the private value of experimentation need not align across partners (a problem that does not emerge in standard auctions for physical goods). For example, after a few interactions, agent i from side A may have learned his value for agent j from side B , while agent j may face residual uncertainty about his value for agent i . The scores $S^{I;\beta}$ in the proposed auctions are thus different from a simple combination of the Gittins indexes corresponding to the agents' own private values for experimentation. They are constructed to internalize the cost and benefits of cross-subsidization, as perceived by the platform, while also accounting for the costs of the agents' private information.

One of the appeals of the proposed auctions is that they admit simple equilibria in which agents bid myopically their match values in each period. In other words, the agents do not need to know how to solve complex dynamic-programming problems, or be able to compute the indexes (although, nowadays, there is software that does so). Once the scoring and payments rules are understood, it is in the agents' interest to bid "straight-forwardly" in all periods, irrespective of the beliefs they may have about other agents' past and current types, and irrespective of their own, as well as other agents', past behavior. As in other market design settings, however, it is important that the market designer educates the bidders by carefully

explaining the structure of the scoring and payment rules and why the proposed auctions admit such simple equilibria. Also note that, while the proposed auctions are, in spirit, similar to the familiar VCG auctions used in many markets, they are adapted to control for the dynamic nature of the externalities the agents impose on one another, and for the fact that the matches implemented in equilibrium need not maximize total surplus.

7 Related Literature

Dynamics in our model stem from changes in actual, or perceived, match values. To the best of our knowledge, this is the first paper to study dynamic matching in environments in which agents receive multi-dimensional new private information in each period, revise their match values, and change partners multiple times. The analysis is related to the following strands of the literature.

Position and Scoring Auctions. In procurement auctions, scoring rules are often used to aggregate the various dimensions of the sellers' offers (price, product design, delivery time, etc.). See, for example, Che (1993), and Asker and Cantillon (2008). Our matching auctions share with this literature the idea that the desired allocations can be induced through an appropriate design of the scoring rules governing the auctions. However, while the above literature focuses on static settings, our scoring rules are for dynamic environments in which preferences evolve over time. More importantly, the scores in our auctions aggregate the preferences of different agents from different sides of the market, whereas in the procurement auctions they aggregate the various dimensions of each seller's own offer.

Another related literature studies auctions for sponsored links by search engines and for contextualized ads by ad exchanges. For example, Varian (2007), Edelman, Ostrovsky and Schwarz (2007), and Gomes and Sweeney (2014) study the properties of the GSP auction used by online search engines to allocate ads, while Mansour, Muthukrishnan, and Nisan (2012), Harris and Varian (2014), and Arnosti, Beck and Milgrom (2016) study auctions used by ad exchanges to match advertisers with content providers.³¹ Our matching auctions are relevant also for these markets, modulo the fact that, in online search, searchers typically do not pay for matches (or, more precisely, the currency used for the services they receive is the release of their privacy). In ad exchange auctions, instead, it is becoming customary for content providers to specify dynamic reservation prices for different types of ads, which is a form of dynamic bidding. In these auctions, both sides thus repeatedly bid for all matches, bids are aggregated into bilateral scores, and participating bidders are charged upfront fees at the time of joining, as in our model.

³¹See also Athey and Ellison (2011), Börgers, Cox, Pesendorfer and Petricek (2013), and Gomes (2014).

Platform Markets. Markets where agents purchase access to other agents are the focus of a vast literature on two-sided markets (see Belleflamme and Peitz (2017) for a recent overview). This literature restricts attention to a single network, or to mutually exclusive networks. Most importantly, it focuses primarily on static environments (see Cabral (2011) for a dynamic model with complete information, and Jullien and Pavan (2018) for a dynamic extension with asymmetric information; in these works, however, match values are constant over time).

Dynamic Matching. Most of the recent literature on centralized dynamic matching focuses on markets without transfers, in which matching is irreversible and dynamics originate in the arrival and departure of agents to and from the market. In the context of kidney exchange, for example, Ünver (2010) studies optimal mechanisms for two-way and multi-way exchanges, minimizing total waiting costs. Optimal dynamic matching is also the focus of Anderson, Ashlagi, Gamarnik, and Kanoria (2015), Baccara, Lee, and Yariv (2015), Akbarpour, Li, and Oveis Gharan (2016), and Herbst and Schickner (2016). A key trade-off in such environments is between avoiding waiting costs and waiting for the market to thicken. A related strand of papers studies the assignment of objects through waiting lists (see Leshno (2015), Thakral (2015), Bloch and Cantala (2016), and Schummer (2016) for recent developments).³² The key differences with respect to the present paper are that, in this literature, match values are constant over time and platforms do not use payments to incentivize the agents.

Matching with Transfers. Damiano and Li (2007) and Johnson (2013) consider the mechanism-design problem of a profit-maximizing platform facing agents with private information on their vertical types. Hoppe, Moldovanu and Ozdenoren (2011) quantify the benefit of a coarse matching scheme in terms of matching surplus, revenue, and welfare. In these papers, in equilibrium, matching is one-to-one. In contrast, Board (2009) considers the problem of a profit-maximizing platform allocating agents to mutually exclusive groups (e.g., teams), whereas Gomes and Pavan (2016, 2018) study a problem similar to the one in Board (2009) but where agents differ both in their preferences and in their attractiveness and matching is non-partitional. Hoppe, Moldovanu and Sela (2009) show that assortative matching can arise in a Bayesian equilibrium of a bilateral (costly) signaling game. Dizdar and Moldovanu (2016) study a model where agents are characterized by private, multi-dimensional, attributes which jointly determine the surplus from a match, and give a possible explanation for the prevalence of rules that divide surplus in a fixed proportion. Matching in all of these papers is static. In contrast, Fershtman and Pavan (2017) consider a simple model of dynamic matching which is a special version of the model in Section 6 in which valuations change only once, after the first

³²See also Damiano and Lam (2005), Kurino (2009), and Doval (2015) for appropriate stability concepts in dynamic matching environments.

interaction, a single match is formed in each period, and the new private information agents receive in each period is uni-dimensional.

Dynamic Mechanism Design. From a methodological standpoint, we draw from recent developments in the dynamic mechanism design literature. In particular, the conditions for incentive compatibility in the present paper adapt to the matching environment under examination results in Bergemann and Valimaki (2010) and Pavan, Segal, and Toikka (2014). See also Baron and Besanko (1984), Besanko (1985), Courty and Li (2000), Board (2007), and Eso and Szentes (2007), for some of the earlier contributions, and Gershkov and Moldovanu (2014), Borgers (2015), Bergemann and Pavan (2015), Bergemann and Valimaki (2017), and Pavan (2017) for overviews of this literature. As explained above, the optimality of truthful (myopic) bidding in periods other than the initial one in the present paper follows from arguments similar to those in Bergemann and Valimaki (2010) and Kakade et al. (2013). The optimality of joining the auctions in period zero and selecting the membership status designed for each vertical type, instead, follows from the two monotonicity conditions discussed in the previous sections. These conditions are similar in spirit to the integral-monotonicity conditions in Pavan, Segal, and Toikka (2014), but are adapted to account for the multi-dimensionality of the private information the agents receive in each period, after joining the platform.

8 Conclusions

This paper introduces and then studies a class of dynamic matching auctions for markets in which agents' preferences for potential partners evolve over time.

The proposed auctions are fairly simple to operate and can be used in a variety of markets including those for scientific outsourcing, peer-to-peer lending, ad exchanges, and organized events. Upon joining the platform, agents are asked to select a membership status which determines the weight assigned to their bids in the subsequent auctions. They then bid repeatedly for potential partners from the opposite side of the market. In each period, the platform computes bilateral scores, one for each match, and implements those matches that maximize the sum of the scores, subject to individual and aggregate capacity constraints. We show that, under reasonable conditions, such a class includes auctions that maximize the platform's profits as well as auctions that maximize total welfare, over all possible mechanisms. The analysis also sheds light on the inefficiencies due to profit maximization, which in turn can be used to guide government interventions in certain markets. Finally, we show how similar auctions but with forward-looking scores taking the form of indexes can be used in certain markets in which the evolution of the match values is endogenous, possibly due to experimentation, a preference for variety, of habit formation.

In future work, it would be interesting to study how the auctions must be adapted if one side does not bid for the matches, as is currently the case in sponsored search auctions. It would also be interesting to consider markets in which the platform must incur a cost only when it matches a pair who was not matched in the preceding period. Such costs introduce additional non-separabilities in the matching allocations that translate into a certain form of inertia and give rise to a trade-off between improving the quality of the existing matches and economizing on future re-matching costs, which is absent in the present analysis.

Another fruitful direction for future research consists in comparing the matching dynamics in centralized markets such as those investigated in the present paper to their counterparts in decentralized markets. To the best of our knowledge, there is no tractable model of decentralized matching where agents perform on-the-job search and rematch over time. Developing such a model is challenging but is an important next step for this literature.

Matching dynamics in the present paper originate in changes in agents' preferences for potential partners. Another strand of the matching literature explores matching dynamics driven by the (stochastic) arrival or departure of agents to and from the market (see, for example, Anderson et al. 2015, Baccara et al. 2015, Akbarpour et al. 2016). Combining stochastic arrivals and departures with time-varying match values is another interesting direction for future research.

Appendix

Proof of Theorem 1. The proof is in two steps. Step 1 shows that, fixing the membership statuses, when the period- t payments are as in (5), for any $t \geq 1$, any period- $(t-1)$ history, participating in the auctions and following truthful strategies constitutes a *periodic ex-post continuation equilibrium* for the game starting with the period- $(t-1)$ history. Step 2 first shows that the matches implemented under truthful bidding satisfy certain monotonicity conditions (defined below). It then shows that such monotonicities imply that, when the membership fees are as in (6), then participating in period zero and selecting a membership status equal to the true vertical type is optimal for any individual who expects all other individuals to follow truthful strategies, irrespective of the individual's beliefs about the opponents' vertical types.

Step 1. Fix the weights $\hat{\beta} \equiv \beta(\hat{\theta})$ determined by the agents' selection of the membership statuses $\hat{\theta}$ at $t = 0$. Denote by $\hat{\chi}$ an arbitrary matching rule that, at each period $t \geq 1$, given the weights $\hat{\beta}$ and the received period- t bids b_t , implements the matches for which the sum of the scores $\hat{S}_{ijt}(b_t) \equiv \hat{\beta}_i^A b_{ijt}^A + \hat{\beta}_j^B b_{ijt}^B - c_{ijt}$ is the highest, subject to aggregate and individual capacity constraints. Similarly, let $\hat{\chi}^{-l,k}$ be any matching rule that maximizes the sum of the scores subject to aggregate and individual capacity constraints, when the score of

any match that involves agent l from side k is identically equal to zero (in this case, the rule $\hat{\chi}^{-l,k}$ can be assumed to never implement any match involving agent l from side k). Denote by $\hat{\psi}_{t \geq 1} \equiv (\hat{\psi}_s)_{s \geq 1}$ the collection of payment functions given by (5), when the selection of the membership statuses is $\hat{\theta}$. Finally, denote by \hat{w}_t and by $\hat{w}_t^{-l,k}$ the flow period- t total weighted surplus implemented under the rules $\hat{\chi}$ and $\hat{\chi}^{-l,k}$, respectively, when the weights are given by $\hat{\beta}$.

Lemma 1 (Optimality for $t \geq 1$). *Fix the selection $\hat{\theta}$ of the membership statuses at $t = 0$. Starting from any period- $(t - 1)$ history, participating in the auction and following truthful strategies constitutes a periodic ex-post continuation equilibrium for the continuation game that starts in period t with the period- $(t - 1)$ history, $t \geq 1$.*

Proof of Lemma 1. We show that, in the continuation game that starts in period $t \geq 1$, irrespective of the history of past play, of the true vertical type profile θ , and of the history of past and current horizontal types $\varepsilon^t \equiv (\varepsilon_s)_{s=1}^t$, any agent $l \in N^k$, $k = A, B$, who expects all other agents to participate and follow truthful strategies from period t (included) onwards, finds it optimal to do the same.

Consider agent l from side A (the problem for any side- B agent is similar). Suppose that the true profile of vertical types is θ , and the true profile of period- t horizontal types is ε_t . Recall that, in the proposed auctions, the matches and the payments at each period $s > t$ are invariant to the period- t bids. To prove the result, it thus suffices to show that the flow surplus the agent obtains in period t by bidding truthfully is higher than the flow surplus he obtains by submitting any other vector of bids. In other words, it suffices to show that, for all possible bids $b_{lt}^A = (b_{ljt}^A)_{j \in N^B}$ of agent l

$$\sum_{j \in N^B} v_{ljt}^A \hat{\chi}_{ljt} \left(v_{lt}^A, v_t^{-l,A} \right) - \hat{\psi}_{lt}^A \left(v_{lt}^A, v_t^{-l,A} \right) \geq \sum_{j \in N^B} v_{ljt}^A \hat{\chi}_{ljt} \left(b_{lt}^A, v_t^{-l,A} \right) - \hat{\psi}_{lt}^A \left(b_{lt}^A, v_t^{-l,A} \right), \quad (16)$$

where $v_{lt}^A = \theta_{lt}^A \varepsilon_{lt}^A$. When the payments are as in (5), the left hand side of the above inequality is equal to

$$\frac{1}{\hat{\beta}_l^A} \left[\hat{w}_t \left(v_{lt}^A, v_t^{-l,A} \right) - \hat{w}_t^{-l,A} \left(v_{lt}^A, v_t^{-l,A} \right) \right], \quad (17)$$

where $\hat{w}_t \left(v_{lt}^A, v_t^{-l,A} \right) \equiv \sum_{i \in N^A} \sum_{j \in N^B} \hat{S}_{ijt} \left(v_{lt}^A, v_t^{-l,A} \right) \hat{\chi}_{ijt} \left(v_{lt}^A, v_t^{-l,A} \right)$.

Furthermore, the period- t payment in the right-hand side of (16) is equal to

$$\begin{aligned}\hat{\psi}_{lt}^A \left(b_{lt}^A, v_t^{-l,A} \right) &= \sum_{j \in N^B} b_{ljt}^A \hat{\chi}_{ljt} \left(b_{lt}^A, v_t^{-l,A} \right) - \frac{1}{\hat{\beta}_t^A} \left[\hat{w}_t \left(b_{lt}^A, v_t^{-l,A} \right) - \hat{w}_t^{-l,A} \left(b_{lt}^A, v_t^{-l,A} \right) \right] \\ &= -\frac{1}{\hat{\beta}_t^A} \sum_{i \in N^A \setminus \{l\}} \sum_{j \in N^B} \hat{S}_{ijt} \left(b_{lt}^A, v_t^{-l,A} \right) \cdot \hat{\chi}_{ijt} \left(b_{lt}^A, v_t^{-l,A} \right) \\ &\quad - \frac{1}{\hat{\beta}_t^A} \sum_{j \in N^B} \left(\hat{\beta}_j^B v_{ljt}^B - c_{ljt} \right) \cdot \hat{\chi}_{ljt} \left(b_{lt}^A, v_t^{-l,A} \right) + \frac{1}{\hat{\beta}_t^A} \hat{w}_t^{-l,A} \left(b_{lt}^A, v_t^{-l,A} \right).\end{aligned}$$

This means that the right hand side of the inequality in (16) is equal to

$$\frac{1}{\hat{\beta}_t^A} \sum_{i \in N^A} \sum_{j \in N^B} \hat{S}_{ijt} \left(v_{lt}^A, v_t^{-l,A} \right) \cdot \hat{\chi}_{ijt} \left(b_{lt}^A, v_t^{-l,A} \right) - \frac{1}{\hat{\beta}_t^A} \hat{w}_t^{-l,A} \left(b_{lt}^A, v_t^{-l,A} \right).$$

Next, observe that $\hat{w}_t^{-l,A}$ is invariant to the bids by agent l from side k , meaning that

$$\hat{w}_t^{-l,A} \left(v_{lt}^A, v_t^{-l,A} \right) = \hat{w}_t^{-l,A} \left(b_{lt}^A, v_t^{-l,A} \right).$$

It follows that the inequality in (16) holds if and only if

$$\hat{w}_t \left(v_{lt}^A, v_t^{-l,A} \right) \geq \sum_{i \in N^A} \sum_{j \in N^B} \hat{S}_{ijt} \left(v_{lt}^A, v_t^{-l,A} \right) \cdot \hat{\chi}_{ijt} \left(b_{lt}^A, v_t^{-l,A} \right). \quad (18)$$

The inequality in (18) follows from the definition of the matching rule $\hat{\chi}$. This means that bidding truthfully yields the agent a payoff at least as high as under any other bidding strategy. Finally, to see that it is also (periodically ex-post) optimal for the agent to participate at all periods $t \geq 1$, after any history, note that the agent's continuation payoff under truthful strategies is proportional to his expected contribution to continuation weighted surplus, which is always nonnegative given that $\hat{\beta}_t^A > 0$ and that, at each period, $\hat{w}_t \left(v_{lt}^A, v_t^{-l,A} \right) - \hat{w}_t^{-l,A} \left(v_{lt}^A, v_t^{-l,A} \right) \geq 0$. ■

Step 2. We now show that, when the period-0 membership fees are as in (6), participating in period zero and selecting a membership status equal to the true vertical type is optimal for any individual who expects all other individuals to participate and follow truthful strategies, irrespective of the individual's beliefs about the opponents' vertical types.

Let

$$D_l^k(\theta) \equiv \begin{cases} \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{h \in N^B} \varepsilon_{lht}^A \chi_{lht} \mid \theta \right] & \text{if } k = A \\ \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{h \in N^A} \varepsilon_{hlt}^B \chi_{hlt} \mid \theta \right] & \text{if } k = B \end{cases} \quad (19)$$

denote the “quality” of the matches expected by agent $l \in N^k$ when the profile of the agents’ true vertical types is θ and all agents are expected to follow truthful strategies. Similarly, let

$$\tilde{D}_l^k(\hat{\theta}_l^k, (\theta_l^k, \theta^{-l,k})) \equiv \begin{cases} \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{h \in N^B} \varepsilon_{lht}^A \chi_{lht}(\hat{\theta}_l^A, \theta^{-l,A}, b_t) \mid \theta \right] & \text{if } k = A \\ \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{h \in N^A} \varepsilon_{hlt}^B \chi_{hlt}(\hat{\theta}_l^B, \theta^{-l,B}, b_t) \mid \theta \right] & \text{if } k = B \end{cases},$$

denote the match quality that agent $l \in N^k$ expects when the true profile of vertical types is $\theta = (\theta_l^k, \theta^{-l,k}) \in \Theta$, the agent selects the membership status $\hat{\theta}_l^k$ at $t = 0$ and then bids truthfully from $t = 1$ onwards, and all agents other than l follow truthful strategies at each period, including period zero. Note that the reason why we distinguish between the case in which $k = A$ and the one in which $k = B$ in the expressions for $D_l^k(\theta)$ and $\tilde{D}_l^k(\hat{\theta}_l^k, (\theta_l^k, \theta^{-l,k}))$ is that the order in the subscripts of the allocations χ_{ijt} , as well as the order in the subscripts in the horizontal types ε_{ijt}^k , is not permutable: the first index always refers to side A , while the second to side B . Also note that the expectation is over the bids, given the true vertical types θ . Finally, note that, when $\hat{\theta}_l^k = \theta_l^k$, for any θ ,

$$\tilde{D}_l^k(\theta_l^k, (\theta_l^k, \theta^{-l,k})) = D_l^k(\theta). \quad (20)$$

Indeed, the only difference between the two functions is that \tilde{D}_l^k accommodates for the possibility that agent l ’s membership status differs from his true vertical type.

For any agent $i \in N^A$ (the arguments for the side- B agents are analogous), then let

$$\hat{U}_i^A(\theta) \equiv \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j \in N^B} \theta_i^A \varepsilon_{ijt}^A \chi_{ijt}(\theta, b_t) \mid \theta \right] - \mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t \psi_{it}^A(\theta, b_t) \mid \theta \right]$$

denote the payoff that the agent expects in the matching auctions when the true vertical type profile is θ and all agents follow truthful strategies at all periods. Note that

$$\mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t \psi_{it}^A(\theta, b_t) \mid \theta \right] = \theta_l^k D_l^k(\theta) - \int_{\underline{\theta}_l^k}^{\theta_l^k} D_l^k(y, \theta^{-l,k}) dy - L_l^k,$$

which implies that

$$\hat{U}_i^k(\theta) = \int_{\hat{\theta}_i^k}^{\theta_i^k} D_i^k(y, \theta^{-l,k}) dy + L_i^k. \quad (21)$$

The next lemma, which is the key step in the proof, shows that the quality of the matches expected under truthful bidding is always nondecreasing in the selected membership status, both when the latter coincides with the agent's true vertical type (part (ii) in the lemma) as well as when it does not (part (i) in the lemma). The first property guarantees the monotonicity of the agents' equilibrium payoffs in their vertical type, a property we use to establish Theorem 2, whereas the second property guarantees the optimality of selecting a membership status equal to the true vertical type at $t = 0$.

Lemma 2 ((Key monotonicities)). *For all $l \in N^k$, $k = A, B$, the following monotonicities hold:*

- (i) $\tilde{D}_i^k(\hat{\theta}_i^k, \theta)$ is non-decreasing in $\hat{\theta}_i^k$, all $\theta \in \Theta$;
- (ii) $D_i^k(\theta_i^k, \theta^{-l,k})$ is non-decreasing in θ_i^k , all $\theta^{-l,k} \in \Theta^{-l,k}$.

Proof of Lemma 2. Consider an arbitrary agent $i \in N^A$ from side A (the arguments for the side- B agents are analogous) and fix the profile of types $\theta^{-i,A}$ for the other agents. We establish part (ii) first.

Part (ii). Take any pair of types $\theta_i^A, \hat{\theta}_i^A \in \Theta_i^A$, with $\theta_i^A < \hat{\theta}_i^A$, and let $\chi_t(\theta, b_t)$ denote the period- t matches implemented under the proposed auctions when the selected membership statuses are θ and the period- t bids are b_t . That, at each period, the auctions select the matches that maximize the sum of the scores (with the weights determined by the membership statuses), subject to the aggregate and the individual capacity constraints, implies that

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{r \in N^A \setminus \{i\}} \sum_{j \in N^B} (\beta_r^A(\theta_r^A) b_{rjt}^A + \beta_j^B(\theta_j^B) b_{rjt}^B - c_{rjt}) \chi_{rjt} \left((\hat{\theta}_i^A, \theta^{-i,A}), \left((\hat{\theta}_i^A \varepsilon_{ijt}^A)_{j \in N^B}, b_t^{-i,A} \right) \right) \mid (\hat{\theta}_i^A, \theta^{-i,A}) \right] \\ & + \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j \in N^B} \left(\beta_i^A(\hat{\theta}_i^A) (\hat{\theta}_i^A \varepsilon_{ijt}^A) + \beta_j^B(\theta_j^B) b_{ijt}^B - c_{ijt} \right) \chi_{ijt} \left((\hat{\theta}_i^A, \theta^{-i,A}), \left((\hat{\theta}_i^A \varepsilon_{ijt}^A)_{j \in N^B}, b_t^{-i,A} \right) \right) \mid (\hat{\theta}_i^A, \theta^{-i,A}) \right] \\ & \geq \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{r \in N^A \setminus \{i\}} \sum_{j \in N^B} (\beta_r^A(\theta_r^A) b_{rjt}^A + \beta_r^B(\theta_r^B) b_{rjt}^B - c_{rjt}) \chi_{rjt} \left((\theta_i^A, \theta^{-i,A}), \left((\theta_i^A \varepsilon_{ijt}^A)_{j \in N^B}, b_t^{-i,A} \right) \right) \mid (\hat{\theta}_i^A, \theta^{-i,A}) \right] \\ & + \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j \in N^B} \left(\beta_i^A(\hat{\theta}_i^A) (\hat{\theta}_i^A \varepsilon_{ijt}^A) + \beta_j^B(\theta_j^B) b_{ijt}^B - c_{ijt} \right) \chi_{ijt} \left((\theta_i^A, \theta^{-i,A}), \left((\theta_i^A \varepsilon_{ijt}^A)_{j \in N^B}, b_t^{-i,A} \right) \right) \mid (\hat{\theta}_i^A, \theta^{-i,A}) \right] \\ & = \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{r \in N^A \setminus \{i\}} \sum_{j \in N^B} (\beta_r^A(\theta_r^A) b_{rjt}^A + \beta_r^B(\theta_r^B) b_{rjt}^B - c_{rjt}) \chi_{rjt} \left((\theta_i^A, \theta^{-i,A}), \left((\theta_i^A \varepsilon_{ijt}^A)_{j \in N^B}, b_t^{-i,A} \right) \right) \mid (\theta_i^A, \theta^{-i,A}) \right] \\ & + \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j \in N^B} \left(\beta_i^A(\hat{\theta}_i^A) (\hat{\theta}_i^A \varepsilon_{ijt}^A) + \beta_j^B(\theta_j^B) b_{ijt}^B - c_{ijt} \right) \chi_{ijt} \left((\theta_i^A, \theta^{-i,A}), \left((\theta_i^A \varepsilon_{ijt}^A)_{j \in N^B}, b_t^{-i,A} \right) \right) \mid (\theta_i^A, \theta^{-i,A}) \right]. \end{aligned}$$

The left-hand side of the above inequality is the expected weighted surplus — under the weights $(\beta_i^A(\hat{\theta}_i^A), \beta^{-i,A}(\theta^{-i,A}))$ — when all agents follow truthful strategies from period $t = 0$ onward and the true profile of vertical types is $(\hat{\theta}_i^A, \theta^{-i,A})$. The right-hand side of the inequality is the expected weighted surplus — under the same weights $(\beta_i^A(\hat{\theta}_i^A), \beta^{-i,A}(\theta^{-i,A}))$ — when, at the same profile of true vertical types $(\hat{\theta}_i^A, \theta^{-i,A})$, all agents other than agent i from side A follow truthful strategies in all periods whereas agent i from side A perfectly replicates the behavior of type θ_i^A in each period (that is, he selects the membership status θ_i^A at $t = 0$ and then, at each subsequent period, given the true horizontal types $\varepsilon_{it}^A \equiv (\varepsilon_{ijt}^A)_{j=1,\dots,n_B}$, submits bids equal to $b_{ijt}^A = \theta_i^A \varepsilon_{ijt}^A$, all $j \in N^B$). The expectations are with respect to the horizontal types given the vertical types. The inequality follows from the fact that, holding the weights $(\beta_i^A(\hat{\theta}_i^A), \beta^{-i,A}(\theta^{-i,A}))$ fixed in the computation of the flow surpluses, at each period, given the bids $\left(\left(\hat{\theta}_i^A \varepsilon_{ijt}^A \right)_{j \in N^B}, b_t^{-i,A} \right)$, the matches $\chi_t \left((\hat{\theta}_i^A, \theta^{-i,A}), \left(\left(\hat{\theta}_i^A \varepsilon_{ijt}^A \right)_{j \in N^B}, b_t^{-i,A} \right) \right)$ implemented when the membership status selected by agent i leads to the same weight $\beta_i^A(\hat{\theta}_i^A)$ used to compute the flow surpluses maximize the flow surplus. The equality in the above expression then follows from the independence of the horizontal types from the vertical ones.

The above inequality is not to be confused with agent i 's incentive-compatibility constraints. As explained in the main text, the monotonicity of match quality in the lemma does not follow from standard arguments in (both static and dynamic) screening models where the monotonicity of the allocations follows from the combination of incentive compatibility with the supermodularity in the agents' payoffs (between vertical types and allocations).

Now, inverting the role of $\hat{\theta}_i^A$ and θ_i^A in the above inequality, we have that

$$\begin{aligned}
& \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{r \in N^A \setminus \{i\}} \sum_{j \in N^B} (\beta_r^A(\theta_r^A) b_{rjt}^A + \beta_j^B(\theta_j^B) b_{rjt}^B - c_{rjt}) \chi_{rjt} \left((\theta_i^A, \theta^{-i,A}), \left((\theta_i^A \varepsilon_{ijt}^A)_{j \in N^B}, b_t^{-i,A} \right) \right) \mid (\theta_i^A, \theta^{-i,A}) \right] \\
& + \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j \in N^B} (\beta_i^A(\theta_i^A) (\theta_i^A \varepsilon_{ijt}^A) + \beta_j^B(\theta_j^B) b_{ijt}^B - c_{ijt}) \chi_{ijt} \left((\theta_i^A, \theta^{-i,A}), \left((\theta_i^A \varepsilon_{ijt}^A)_{j \in N^B}, b_t^{-i,A} \right) \right) \mid (\theta_i^A, \theta^{-i,A}) \right] \\
& \geq \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{r \in N^A \setminus \{i\}} \sum_{j \in N^B} (\beta_r^A(\theta_r^A) b_{rjt}^A + \beta_r^B(\theta_r^B) b_{rjt}^B - c_{rjt}) \chi_{rjt} \left((\hat{\theta}_i^A, \theta^{-i,A}), \left((\hat{\theta}_i^A \varepsilon_{ijt}^A)_{j \in N^B}, b_t^{-i,A} \right) \right) \mid (\hat{\theta}_i^A, \theta^{-i,A}) \right] \\
& + \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j \in N^B} (\beta_i^A(\theta_i^A) (\theta_i^A \varepsilon_{ijt}^A) + \beta_j^B(\theta_j^B) b_{ijt}^B - c_{ijt}) \chi_{ijt} \left((\hat{\theta}_i^A, \theta^{-i,A}), \left((\hat{\theta}_i^A \varepsilon_{ijt}^A)_{j \in N^B}, b_t^{-i,A} \right) \right) \mid (\hat{\theta}_i^A, \theta^{-i,A}) \right].
\end{aligned}$$

Combining the results in the above two inequalities, we have that

$$\left(\beta_i^A(\hat{\theta}_i^A) \hat{\theta}_i^A - \beta_i^A(\theta_i^A) \theta_i^A \right) \cdot \left(D_i^A((\hat{\theta}_i^A, \theta^{-i,A})) - D_i^A((\theta_i^A, \theta^{-i,A})) \right) \geq 0.$$

Because $\beta_i^A(\cdot)$ is strictly positive and non-decreasing and $\hat{\theta}_i^A > \theta_i^A$, it must be that

$$D_i^A((\hat{\theta}_i^A, \theta^{-i,A})) \geq D_i^A(\theta_i^A, \theta^{-i,A}).$$

Part (i). Again, because, in each period, the matches implemented under truthful strategies maximize the sum of the scores subject to the aggregate and the individual capacity constraints, we have that, for any $\hat{\theta}_i^A, \theta_i^A \in \Theta_i^A$, $\theta \in \Theta$,

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{r \in N^A \setminus \{i\}} \sum_{j \in N^B} (\beta_r^A(\theta_r^A) b_{rjt}^A + \beta_j^B(\theta_j^B) b_{rjt}^B - c_{rjt}) \chi_{rjt} \left((\hat{\theta}_i^A, \theta^{-i,A}), \left((\theta_i^A \varepsilon_{ijt}^A)_{j \in N^B}, b_t^{-i,A} \right) \right) \mid (\theta_i^A, \theta^{-i,A}) \right] \\ & + \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j \in N^B} \left(\beta_i^A(\hat{\theta}_i^A) (\theta_i^A \varepsilon_{ijt}^A) + \beta_j^B(\theta_j^B) b_{ijt}^B - c_{ijt} \right) \chi_{ijt} \left((\hat{\theta}_i^A, \theta^{-i,A}), \left((\theta_i^A \varepsilon_{ijt}^A)_{j \in N^B}, b_t^{-i,A} \right) \right) \mid (\theta_i^A, \theta^{-i,A}) \right] \\ & \geq \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{r \in N^A \setminus \{i\}} \sum_{j \in N^B} (\beta_r^A(\theta_r^A) b_{rjt}^A + \beta_r^B(\theta_r^B) b_{rjt}^B - c_{rjt}) \chi_{rjt} \left((\theta_i^A, \theta^{-i,A}), \left((\theta_i^A \varepsilon_{ijt}^A)_{j \in N^B}, b_t^{-i,A} \right) \right) \mid (\theta_i^A, \theta^{-i,A}) \right] \\ & + \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j \in N^B} \left(\beta_i^A(\hat{\theta}_i^A) (\theta_i^A \varepsilon_{ijt}^A) + \beta_j^B(\theta_j^B) b_{ijt}^B - c_{ijt} \right) \chi_{ijt} \left((\theta_i^A, \theta^{-i,A}), \left((\theta_i^A \varepsilon_{ijt}^A)_{j \in N^B}, b_t^{-i,A} \right) \right) \mid (\theta_i^A, \theta^{-i,A}) \right]. \end{aligned}$$

The left-hand side of the above inequality is the expected weighted surplus when the true vertical types are $(\theta_i^A, \theta^{-i,A})$, the weights used to compute the flow surpluses are given by $(\beta_i^A(\hat{\theta}_i^A), \beta^{-i,A}(\theta^{-i,A}))$, all agents other than agent i from side A follow truthful strategies, and agent i selects the membership status $\hat{\theta}_i^A$ and then bids truthfully at all periods. The right-hand side of the above inequality is the expected weighted surplus under the same true vertical types $(\theta_i^A, \theta^{-i,A})$, when the weights used to compute the flow surpluses continue to be given by $(\beta_i^A(\hat{\theta}_i^A), \beta^{-i,A}(\theta^{-i,A}))$, and all agents, including agent i from side A , follow truthful strategies at all periods. Once again, the inequality follows from the fact that, when the weights are $(\beta_i^A(\hat{\theta}_i^A), \beta^{-i,A}(\theta^{-i,A}))$, in each period, surplus is higher when the membership status selected by agent i leads to scores containing the same weight $\beta_i^A(\hat{\theta}_i^A)$ used to evaluate the flow surpluses.

By the same arguments,

$$\begin{aligned}
& \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{r \in N^A \setminus \{i\}} \sum_{j \in N^B} (\beta_r^A(\theta_r^A) b_{rjt}^A + \beta_j^B(\theta_j^B) b_{rjt}^B - c_{rjt}) \chi_{rjt} \left((\theta_i^A, \theta^{-i,A}), \left((\theta_i^A \varepsilon_{ijt}^A)_{j \in N^B}, b_t^{-i,A} \right) \right) \mid (\theta_i^A, \theta^{-i,A}) \right] \\
& + \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j \in N^B} (\beta_i^A(\theta_i^A) (\theta_i^A \varepsilon_{ijt}^A) + \beta_j^B(\theta_j^B) b_{ijt}^B - c_{ijt}) \chi_{ijt} \left((\theta_i^A, \theta^{-i,A}), \left((\theta_i^A \varepsilon_{ijt}^A)_{j \in N^B}, b_t^{-i,A} \right) \right) \mid (\theta_i^A, \theta^{-i,A}) \right] \\
\geq & \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{r \in N^A \setminus \{i\}} \sum_{j \in N^B} (\beta_r^A(\theta_r^A) b_{rjt}^A + \beta_r^B(\theta_r^B) b_{rjt}^B - c_{rjt}) \chi_{rjt} \left((\hat{\theta}_i^A, \theta^{-i,A}), \left((\theta_i^A \varepsilon_{ijt}^A)_{j \in N^B}, b_t^{-i,A} \right) \right) \mid (\theta_i^A, \theta^{-i,A}) \right] \\
& + \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j \in N^B} (\beta_i^A(\theta_i^A) (\theta_i^A \varepsilon_{ijt}^A) + \beta_j^B(\theta_j^B) b_{ijt}^B - c_{ijt}) \chi_{ijt} \left((\hat{\theta}_i^A, \theta^{-i,A}), \left((\theta_i^A \varepsilon_{ijt}^A)_{j \in N^B}, b_t^{-i,A} \right) \right) \mid (\theta_i^A, \theta^{-i,A}) \right].
\end{aligned}$$

Combining the above two inequalities, we obtain that

$$\left(\beta_i^A(\hat{\theta}_i^A) - \beta_i^A(\theta_i^A) \right) \cdot \theta_i^A \cdot \left(\tilde{D}_i^A(\hat{\theta}_i^A, \theta) - \tilde{D}_i^A(\theta_i^A, \theta) \right) \geq 0.$$

Because $\theta_i^A > 0$ and because $\beta_i^A(\cdot)$ is strictly positive and non-decreasing, it must be that $\tilde{D}_i^A(\cdot, \theta)$ is non-decreasing.³³ ■

We now show that, when status is priced according to the formula in (6), the monotonicities in the previous lemma imply that each agent finds it optimal to select the membership status designed for his true vertical type. Without loss of generality, take an arbitrary agent $i \in N^A$ from side A (the arguments for the side- B agents are analogous). Let

$$\tilde{U}_i^A(\hat{\theta}_i^A; \theta) \equiv \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j \in N^B} \theta_i^A \varepsilon_{ijt}^A \chi_{ijt} \left((\hat{\theta}_i^A, \theta^{-i,A}), b_t \right) \mid \theta \right] - \mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t \psi_{it}^A \left((\hat{\theta}_i^A, \theta^{-i,A}), b_t \right) \mid \theta \right]$$

denote the intertemporal payoff that the agent expects when the true vertical type profile is θ , the agent chooses the membership status designed for type $\hat{\theta}_i^A$, he bids truthfully from period $t = 1$ onwards, and all other agents follow truthful strategies from period $t = 0$ onwards. Note that

$$\tilde{U}_i^A(\theta_i^A; (\theta_i^A, \theta^{-i,A})) = \hat{U}_i^A(\theta_i^A, \theta^{-i,A}), \tag{22}$$

where \hat{U} is as in (21).

Recall from Step 1 that, irrespective of the selection of the membership statuses, remaining in the auctions and bidding truthfully constitutes a periodic ex-post continuation equilibrium starting from any period-1 history (including those reached off path, by deviations in period

³³Note that, because the only influence of membership statuses on match quality is through their impact on the weights β , if $\beta_i^A(\hat{\theta}_i^A) = \beta_i^A(\theta_i^A)$, then $\tilde{D}_i^A(\hat{\theta}_i^A, \theta) = \tilde{D}_i^A(\theta_i^A, \theta)$.

zero). Now consider a fictitious environment where the agent is constrained to select the status $\hat{\theta}_i^A$ at $t = 0$ but is otherwise free to choose any strategy of his choice for the continuation game that starts at $t = 1$. That remaining in the auctions and bidding truthfully constitutes a periodic ex-post continuation equilibrium in the continuation game that starts with period $t = 1$ implies that $\tilde{U}_i^A(\hat{\theta}_i^A; \theta)$ is a value function for the problem the agent faces in the fictitious environment. Standard envelope arguments (see, e.g., Milgrom and Segal (2002)) then imply that $\tilde{U}_i^A(\hat{\theta}_i^A; \theta)$ is Lipschitz continuous in the agent's true vertical type and that it admits the following representation

$$\tilde{U}_i^A(\hat{\theta}_i^A; \theta) = \tilde{U}_i^A(\hat{\theta}_i^A; (\hat{\theta}_i^A, \theta^{-i,A})) + \int_{\hat{\theta}_i^A}^{\theta_i^A} \tilde{D}_i^A(\hat{\theta}_i^A; (y, \theta^{-i,A})) dy. \quad (23)$$

Combining (21) with (23), we then have that

$$\begin{aligned} \tilde{U}_i^A(\hat{\theta}_i^A; \theta) &= \tilde{U}_i^A(\hat{\theta}_i^A; (\hat{\theta}_i^A, \theta^{-i,A})) + \int_{\hat{\theta}_i^A}^{\theta_i^A} \tilde{D}_i^A(\hat{\theta}_i^A; (y, \theta^{-i,A})) dy \\ &\leq \tilde{U}_i^A(\hat{\theta}_i^A; (\hat{\theta}_i^A, \theta^{-i,A})) + \int_{\hat{\theta}_i^A}^{\theta_i^A} \tilde{D}_i^A(y; (y, \theta^{-i,A})) dy \\ &= \hat{U}_i^k(\hat{\theta}_i^A, \theta^{-i,A}) + \int_{\hat{\theta}_i^A}^{\theta_i^A} D_i^A(y, \theta^{-i,A}) dy = \hat{U}_i^k(\theta), \end{aligned}$$

where the first equality follows from (23), the inequality follows from part (i) in Lemma 2, the second equality follows from (20) and (22), and the last equality from (21).

Hence, given θ , the agent prefers following a truthful strategy from period zero onwards than deviating in period zero and then following a truthful strategy from period one onwards.

Finally, because D_i^A is uniformly bounded by E_i^A , participation constraints are satisfied when $L_i^A \geq (\bar{\theta}_i^A - \underline{\theta}_i^A) E_i^A$.

Combining the results in Step 2 with those in Step 1, we thus have that, when L_l^k is large enough, all $l \in N^k$, $k = A, B$, participating in the auctions and following a truthful strategy is a periodic ex-post equilibrium in the entire game. Q.E.D.

Proof of Theorem 2. Consider any feasible mechanism Γ and any BNE σ of the game induced by Γ . Denote by $\hat{\chi} = (\hat{\chi}_t(\theta, \varepsilon^t))_{t=1}^\infty$ and $\hat{\psi} = (\hat{\psi}_t(\theta, \varepsilon^t))_{t=0}^\infty$ the matching and payment rules, as a function of the true state, induced by σ in Γ . Note that we are allowing here for *any* feasible mechanism; that is, the message and signal spaces may be different than those in the matching auctions.

The platform's profits under (Γ, σ) are equal to

$$\mathbb{E} \left[\sum_{k=A,B} \sum_{l \in N^k} \sum_{t=0}^{\infty} \delta^t \hat{\psi}_{lt}^k(\theta, \varepsilon^t) - \sum_{t=1}^{\infty} \delta^t \sum_{i \in N^A} \sum_{j \in N^B} c_{ijt} \hat{\chi}_{ijt}(\theta, \varepsilon^t) \right]. \quad (24)$$

Alternatively, (24) can be rewritten as follows:

$$\mathbb{E} \left[\sum_{t=1}^{\infty} \sum_{i \in N^A} \sum_{j \in N^B} \delta^t ((\theta_i^A \varepsilon_{ijt}^A + \theta_j^B \varepsilon_{ijt}^B - c_{ijt}) \hat{\chi}_{ijt}(\theta, \varepsilon^t)) - \sum_{k=A,B} \sum_{l \in N^k} U_l^k(\theta_l^k) \right], \quad (25)$$

where $U_l^k(\theta_l^k)$ denotes the period-0 interim expected payoff of agent $l \in N^k$ from side $k = A, B$ when his vertical type is θ_l^k , under the equilibrium σ in the mechanism Γ . Note that we denote the interim payoffs by U_l^k to differentiate them from the interim payoff functions \hat{U}_l^k in the auction, whose domain is the entire profile of vertical types.³⁴

The period-0 participation constraints are satisfied if for all $l \in N^k$, $k = A, B$, $\theta_l^k \in \Theta_l^k$, $U_l^k(\theta_l^k) \geq 0$. Following steps similar to those in Pavan, Segal and Toikka (2014, Theorem 1), we can show that the period-0 (interim) expected payoff of each agent $l \in N^k$, $k = A, B$ must satisfy the following envelope condition

$$U_l^k(\theta_l^k) = U_l^k(\underline{\theta}_l^k) + \int_{\underline{\theta}_l^k}^{\theta_l^k} \mathbb{E} [D_l^k(\theta; \hat{\chi})|y] dy, \quad (26)$$

where the expectation is taken over the entire profile of vertical types θ given agent l 's own vertical type. In (26) we emphasize the dependence of the expected match quality

$$D_l^k(\theta; \hat{\chi}) \equiv \begin{cases} \mathbb{E} [\sum_{t=1}^{\infty} \delta^t \sum_{h \in N^B} \varepsilon_{lht}^k \hat{\chi}_{lht}(\theta, \varepsilon^t) \mid \theta] & \text{if } k = A \\ \mathbb{E} [\sum_{t=1}^{\infty} \delta^t \sum_{h \in N^A} \varepsilon_{hlt}^k \hat{\chi}_{hlt}(\theta, \varepsilon^t) \mid \theta] & \text{if } k = B \end{cases}$$

on the matching rule to make clear that such quality varies with (Γ, σ) . The above envelope condition, together with integration by parts, yields the following representation of the

³⁴Recall that \hat{U}_l^k is the payoff agent l expects in the matching auction when the true vertical type profile is θ and all agents follow truthful strategies at all periods.

platform's profits,

$$\mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \sum_{i \in N^A} \sum_{j \in N^B} \left(\left(1 - \frac{1 - F_i^A(\theta_i^A)}{f_i^A(\theta_i^A)\theta_i^A} \right) \theta_i^A \varepsilon_{ijt}^A + \left(1 - \frac{1 - F_j^B(\theta_j^B)}{f_j^B(\theta_j^B)\theta_j^B} \right) \theta_j^B \varepsilon_{ijt}^B - c_{ijt} \right) \hat{\chi}_{ijt}(\theta, \varepsilon^t) \right] - \sum_{k=A,B} \sum_{l \in N^k} U_l^k(\theta_l^k). \quad (27)$$

The first term, which is only a function of the matching rule $\hat{\chi}$, is the *expected dynamic virtual surplus* (DVS) generated by the matching rule $\hat{\chi}$. Clearly, because such a representation applies to the matching rule generated by any BNE of any mechanism, the above representation applies also to the state-contingent matching rules generated by the truthful strategies in the matching auctions.

Now observe that, when the weights are given by $\beta^P \equiv (\beta_l^{k,P}(\cdot))_{l \in N^k, k=A,B}$, the state-contingent matches implemented under the truthful equilibria of the auctions of Theorem 1 maximize DVS over all feasible state-contingent rules $\hat{\chi}$.

Next, let ψ^P be the payment scheme defined by conditions (5) and (6) when the weights are given by β^P and $L_l^k = 0$, all $l \in N^k$, $k = A, B$. From (5) and (6), it is easy to see that, under the equilibria in truthful strategies of the matching auctions where (a) the scoring rule is the one with weights given by β^P , and (b) the payments are ψ^P , the payoff expected by the lowest vertical type of each agent is exactly equal to zero (that is, $U_l^k(\theta_l^k) = 0$, all $l \in N^k$, $k = A, B$). This means that the truthful equilibria of the above matching auctions maximize both terms of (27). Provided all the period-0 participation constraints are satisfied (something we verify below), we then have that the platform's profits are maximized under the truthful equilibria of the proposed matching auctions.

Also note that, while we have restricted attention to (Γ, σ) that yield a deterministic matching rule $\hat{\chi}$, the platform cannot increase its profits by selecting a pair (Γ, σ) that induces a stochastic matching rule. This is because the matches under any such pair (Γ, σ) can also be induced through a deterministic direct mechanism that conditions on the type report of a fictitious agent. The platform's profits under any such a mechanism thus continue to be given by the expression in (27), but with the matching rule conditioning on the behavior of such fictitious agent. This means that the platform's profits are equal to the weighted average of the platform's profits under the deterministic matching rules obtained by conditioning on the various types of the fictitious agent. Because (27) is maximized over all possible deterministic rules under the equilibria in truthful strategies of the proposed matching auctions, we thus have that any pair (Γ, σ) yielding a stochastic matching rule can never improve upon the truthful equilibria of the proposed matching auctions when it comes to the platform's profits.

We now complete the proof by establishing that all period-0 participation constraints are satisfied under the truthful equilibria of the proposed matching auctions. To see this, it suffices to observe that, for all $\theta^{-l,k}$, $l \in N^k$, $k = A, B$, $D_l^k(\theta_l^k, \theta^{-l,k}; \beta^P)$ is non-decreasing in θ_l^k , as established in part (ii) of Lemma 2 in the proof of Theorem 1. The assumption in the theorem that $D_l^k((\theta_l^k, \theta^{-l,k}); \beta^P) \geq 0$, all $l \in N^k$, $k = A, B$, all $\theta^{-l,k}$, then guarantees that $D_l^k(\theta; \beta^P) \geq 0$, all $\theta \in \Theta$, $l \in N^k$, $k = A, B$. Because the period-0 interim payoffs satisfy the envelope condition $\hat{U}_l^k(\theta) = \int_{\theta_l^k}^{\theta_l^k} D_l^k((y, \theta^{-l,k}); \beta^P) dy$ (see the proof of Theorem 1), we then have that $\hat{U}_l^k(\theta) \geq 0$, all $\theta \in \Theta$, $l \in N^k$, $k = A, B$. This means that all the period-0 participation constraints are satisfied (in a periodic ex-post sense, i.e., for any θ , and not just in expectation over $\theta^{-l,k}$ given θ_l^k).³⁵ Q.E.D.

Proof of Theorem 3. As explained in the main text, part (i) follows directly from Theorem 1. Part (ii) follows from arguments similar to those establishing the optimality of the auctions of Theorem 2. In particular, it follows from the fact that (a) the platform's expected profits under any BNE of any mechanism Γ implementing the welfare-maximizing matches satisfy the representation in (27); (b) each agent's period-0 expected payoff satisfies Condition (26); (c) in the proposed auctions, $U_l^k(\underline{\theta}_l^k) = 0$ if, and only if, the payments in (5) and (6) (for $\beta = \beta^W$) are such that $L_l^k = 0$, all $l \in N^k$, $k = A, B$; and (d) when the payments are given by (5) and (6) with $\beta = \beta^W$ and $L_l^k = 0$, all $l \in N^k$, $k = A, B$, all agents' period-0 participation constraints are satisfied, regardless of their beliefs over other agents' types, if, and only if, $D_l^k(\underline{\theta}_l^k, \theta^{-l,k}; \beta^W) \geq 0$. The latter property is a result of the fact that the functions $D_l^k(\cdot, \theta^{-l,k}; \beta^W)$ are non-decreasing, which is shown in part (ii) of Lemma 2 in the proof of Theorem 1. Q.E.D.

Proof of Theorem 4. Let $\hat{\chi}_t^P(\theta, \varepsilon)$ and $\hat{\chi}_t^W(\theta, \varepsilon)$ denote the state-contingent matches implemented in period $t \geq 1$, under the truthful equilibria of, respectively, the profit-maximizing and the welfare-maximizing auctions of Theorems 2 and 3.

Similarly, let $S_{ijt}^P(\theta, \varepsilon)$ and $S_{ijt}^W(\theta, \varepsilon)$ denote the period- t state-contingent scores under the truthful equilibria of the profit-maximizing and the welfare-maximizing auctions of Theorem 2 and 3, respectively. To ease the exposition, hereafter, we drop (θ, ε) from the arguments of all the functions $\hat{\chi}^P$, $\hat{\chi}^W$, S^P , and S^W .

First, observe that, because $\beta_l^{k,P}(\theta_l^k) \leq 1 = \beta_l^{k,W}(\theta_l^k)$, all $\theta_l^k \in \Theta_l^k$, $l \in N^k$, $k = A, B$, when all horizontal types ε are nonnegative, for any $(i, j) \in N^A \times N^B$, any $t \geq 1$, $S_{ijt}^P \leq S_{ijt}^W$.

Part 1. When none of the capacity constraints is binding, in each period $t \geq 1$, the matches implemented under the equilibria of the profit-maximizing auctions (alternatively, the welfare-

³⁵Recall that, under the proposed auctions, $\hat{U}_l^k(\theta_l^k, \theta^{-l,k}) = 0$, all $\theta^{-l,k}$.

maximizing auctions) are all those for which the scores $S_{ijt}^P \geq 0$ (alternatively, $S_{ijt}^W \geq 0$). The above property, along with the fact that, for any $(i, j) \in N^A \times N^B$, $t \geq 1$, $S_{ijt}^W \geq S_{ijt}^P$, then yields the result.

Part 2. Next, consider the case in which only the aggregate capacity constraint is potentially binding, i.e., $M < n^A \cdot n^B$, but $m_l^m \geq n^{-k}$, all $l \in N^k$, $k = A, B$. The result then follows from the following two properties: (a) in each period $t \geq 1$, the set of matches for which $S_{ijt}^P \geq 0$ is a subset of the set of matches for which $S_{ijt}^W \geq 0$, (b) the cardinality of the set of matches implemented in each period in a profit-maximizing auction (alternatively, in a welfare-maximizing auction) is the minimum between M and the cardinality of the set of matches for which $S_{ijt}^P \geq 0$ (alternatively, $S_{ijt}^W \geq 0$).

Part 3. Finally, consider the case in which some of the individual capacity constraints are potentially binding, i.e., $m_l^m < n^{-k}$, for some $l \in N^k$, $k = A, B$. That $S_{ijt}^W \geq S_{ijt}^P$ all $(i, j) \in N^A \times N^B$, all $t \geq 1$, implies that, at any period $t \geq 1$ at which $|\{(i, j) \in N^A \times N^B : \hat{\chi}_{ijt}^P(\theta, \varepsilon) = 1\}| > 0$, necessarily $|\{(i, j) \in N^A \times N^B : \hat{\chi}_{ijt}^W(\theta, \varepsilon) = 1\}| > 0$, for $\hat{\chi}_{ijt}^P(\theta, \varepsilon) = 1$ implies that $S_{ijt}^P(\theta, \varepsilon) \geq 0$ and hence that $S_{ijt}^W(\theta, \varepsilon) \geq 0$. By matching the pair (i, j) the platform then weakly increases welfare relative to a situation in which no pair is matched in period t . Q.E.D.

Proof of Theorem 5. The proof is in 2 parts. Part 1 establishes that conclusions analogous to those in Theorem 1 hold. Part 2 establishes that the same conclusions as in Theorem 4 hold, modulo the restriction to $M = 1$ for the claim in part 2 of Theorem 4. That conclusions analogous to those in Theorem 2 and 3 also hold follows from essentially the same arguments that lead to Theorems 2 and 3, and hence the proof is omitted.

Part 1. The proof is in two steps, whose structure parallels the one in the proof of Theorem 1.

Step 1. First observe that, when all agents follow truthful strategies, the matches implemented in equilibrium maximize continuation weighted surplus, starting from any period- t history, any $t \geq 1$. That is, irrespective of the particular history that led to the selection of the period-0 membership statuses θ_0 and of the past matches x^{t-1} , in the continuation game that starts with period $t \geq 1$, when the true vertical type profile is the one revealed by the selection of the period- t membership statuses, θ_t , the true profile of period- t horizontal types ε_t is the one obtained from the period- t membership statuses θ_t and the period- t bids b_t using (11), the matching rule $\chi^{I;\beta}$ induced by the agents following truthful strategies in the proposed auctions maximizes continuation weighted surplus (13) over the entire set \mathcal{X} of feasible matching rules. This follows from the fact that the problem of maximizing W_t over \mathcal{X} can be seen as a multi-armed bandit problem, with each arm corresponding to a potential

match, and with the flow period- t reward of activating arm (i, j) given by the myopic score $S_{ijt}^{m;\beta} \equiv \beta_i^A(\theta_{i0}^A)v_{ijt}^A + \beta_j^B(\theta_{j0}^B)v_{ijt}^B - c_{ijt}$. When $M = 1$, that $\chi^{I;\beta}$ maximizes W_t follows from known results (see, e.g., Whittle (1982)). When none of the capacity constraints binds (i.e., $M \geq n^A \cdot n^B$, and $m_l^m \geq n^{-k}$, all $l \in N^k$, $k = A, B$), the platform's problem can be viewed as a collection of $n^A \cdot n^B$ separate two-armed bandit problems, one for each potential pair of agents, with the flow reward from matching the pair (i, j) given by $S_{ijt}^{m;\beta}$ and the flow reward from activating the "safe arm" identically equal to zero. It is again well known that the solution to each such problems consists in activating the risky arm if the (i, j) -index $S_{ijt}^{I;\beta}$ is positive and the safe arm if it is negative.

Next, observe that, when the period- t payments are the ones defined in (15), then participating in the auctions and following truthful strategies constitutes a periodic ex-post continuation equilibrium in the continuation game that starts in period $t \geq 1$, irrespective of the period- t history. That remaining in the auctions and following a truthful strategy yields each agent a payoff higher than his outside option follows from the fact that each agent's continuation payoff under truthful strategies is given by $\frac{1}{\beta_l^k(\theta_{l0}^k)}R_{lt}^k$, which is always nonnegative, given that $\beta_l^k(\theta_{l0}^k) > 0$ and that $R_{lt}^k \equiv W_t - W_t^{-l,k}$ is the agent's expected contribution to continuation weighted surplus, which is obviously nonnegative.

To see that, when all other agents follow truthful strategies, then each agent maximizes his expected continuation payoff by following the same strategy, irrespective of his beliefs about other agents' current and past types, consider a representative agent l from side A (the problem for any agent from side B is similar). Suppose that the true profile of vertical types is θ , the true profile of period- t match values is v_t , the profile of period-0 membership choices is θ_0 , and the history of past matches is x^{t-1} .

Denote by $\mathbb{E}^{\tilde{\lambda}[\chi^{I;\beta}]|\theta, v_t, x^{t-1}; (\hat{\theta}_{lt}^A, b_{lt}^A)} [R_{lt+1}^A(\theta_0, \theta_{t+1}, b_{t+1}, x^t)]$ the expected contribution of agent l to continuation weighted surplus from period $t + 1$ onwards, when, in period t , the agent selects the period- t membership status $\hat{\theta}_{lt}^A$, submits the period- t bids b_{lt}^A , follows a truthful strategy from period $t + 1$ onwards, and expects all other agents to follow truthful strategies at all periods $s \geq t$.³⁶ Note that, when the agent follows the truthful strategy also in period t (i.e., when $\hat{\theta}_{lt}^A = \theta_l^A$ and $b_{lt}^A = v_{lt}^A$), then $\tilde{\lambda}[\chi^{I;\beta}]|\theta, v_t, x^{t-1}; (\theta_l^A, v_{lt}^A) = \lambda[\chi^{I;\beta}]|\theta, v_t, x^{t-1}$, where the process $\lambda[\chi^{I;\beta}]|\theta, v_t, x^{t-1}$ is as defined in the main text.

Since the agent can revise his membership status in any of the subsequent periods, any

³⁶The stochastic process $\tilde{\lambda}[\chi^{I;\beta}]|\theta, v_t, x^{t-1}; (\hat{\theta}_{lt}^A, b_{lt}^A)$ is here over future matches, bids and membership choices, under the rule $\chi^{I;\beta}$ that selects in each period the matches maximizing the sum of the index scores, subject to individual and aggregate capacity constraints, when agent l 's period- t choices are $(\hat{\theta}_{lt}^A, b_{lt}^A)$, the true profile of vertical types is θ , the true profile of period- t horizontal types is ε_t , with ε_t obtained from θ and v_t using (11), the history of past matches is x^{t-1} , the agent plans to follow a truthful strategy from $t + 1$ onwards, and all other agents follow truthful strategies from period t onwards.

deviation from the truthful strategy in period t can be corrected in period $t + 1$. This means that, to prove the result, it suffices to show that a single deviation from a truthful strategy is unprofitable to the agent. That is, the agent prefers to follow the truthful strategy from period t onwards than deviating in period t and then reverting to the truthful strategy from period $t + 1$ onwards.

When the agent follows the truthful strategy from period $t + 1$ onwards, his continuation payoff from period $t + 1$ onwards is given by $R_{lt+1}^A(\theta_0, \theta_{t+1}, b_{t+1}, x^t) / \beta_l^A(\theta_{l0}^A)$. Therefore, it is enough to show that, for any period- t selection $(\hat{\theta}_{lt}^A, b_{lt}^A)$,

$$\begin{aligned}
& \sum_{j \in N^B} v_{ljt}^A \chi_{ljt} \left(\theta_0, (\theta_l^A, \theta^{-l,A}), (v_{lt}^A, v_t^{-l,A}), x^{t-1} \right) - \psi_{lt}^A \left(\theta_0, (\theta_l^A, \theta^{-l,A}), (v_{lt}^A, v_t^{-l,A}), x^{t-1} \right) \\
& + \frac{\delta}{\beta_l^A(\theta_{l0}^A)} \mathbb{E}^{\lambda[\chi^{l;\beta}] | \theta_l^A, \theta^{-l,A}, v_{lt}^A, v_t^{-l,A}, x^{t-1}} \left[R_{lt+1}^A(\theta_0, \theta_{t+1}, b_{t+1}, x^t) \right] \\
& \geq \sum_{j \in N^B} v_{ljt}^A \chi_{ljt} \left(\theta_0, (\hat{\theta}_{lt}^A, \theta^{-l,A}), (b_{lt}^A, v_t^{-l,A}), x^{t-1} \right) - \psi_{lt}^A \left(\theta_0, (\hat{\theta}_{lt}^A, \theta^{-l,A}), (b_{lt}^A, v_t^{-l,A}), x^{t-1} \right) \\
& + \frac{\delta}{\beta_l^A(\theta_{l0}^A)} \mathbb{E}^{\tilde{\lambda}[\chi^{l;\beta}] | \theta_l^A, \theta^{-l,A}, v_{lt}^A, v_t^{-l,A}, x^{t-1}; (\hat{\theta}_{lt}^A, b_{lt}^A)} \left[R_{lt+1}^A(\theta_0, \theta_{t+1}, b_{t+1}, x^t) \right]. \tag{28}
\end{aligned}$$

Using the formula for the payments (15) and the definition of the functions R_{lt+1}^A , after some algebra, we have that the left hand side of the above inequality can be rewritten in terms of the functions W_t and $W_t^{-l,A}$ as follows:

$$\frac{1}{\beta_l^A(\theta_{l0}^A)} \left[W_t \left(\theta_0, (\theta_l^A, \theta^{-l,A}), (v_{lt}^A, v_t^{-l,A}), x^{t-1} \right) - W_t^{-l,A} \left(\theta_0, (\theta_l^A, \theta^{-l,A}), (v_{lt}^A, v_t^{-l,A}), x^{t-1} \right) \right].$$

One can also show that the right hand side of the inequality (28) is equal to

$$\begin{aligned}
& \frac{1}{\beta_l^A(\theta_{l0}^A)} \sum_{i \in N^A} \sum_{j \in N^B} S_{ijt}^{m;\beta}(\theta_0, \theta, v_t, x^{t-1}) \cdot \chi_{ijt} \left(\theta_0, (\hat{\theta}_{lt}^A, \theta^{-l,A}), (\hat{b}_{lt}^A, v_t^{-l,A}), x^{t-1} \right) \\
& + \frac{\delta}{\beta_l^A(\theta_{l0}^A)} \mathbb{E}^{\tilde{\lambda}[\chi^{l;\beta}] | \theta_l^A, \theta^{-l,A}, v_{lt}^A, v_t^{-l,A}, x^{t-1}; (\hat{\theta}_{lt}^A, b_{lt}^A)} \left[W_{t+1}(\theta_0, \theta_{t+1}, b_{t+1}, x^t) \right] \\
& - \frac{1}{\beta_l^A(\theta_{l0}^A)} W_t^{-l,A} \left(\theta_0, (\hat{\theta}_{lt}^A, \theta^{-l,A}), (b_{lt}^A, v_t^{-l,A}), x^{t-1} \right).
\end{aligned}$$

Since $W_t^{-l,A}$ is invariant to agent l 's period- t behavior, to establish that the inequality in

(28) holds, it then suffices to show that

$$W_t(\theta_0, \theta, v_t, x^{t-1}) \geq \sum_{i \in N^A} \sum_{j \in N^B} S_{ijt}^{m;\beta}(\theta_0, \theta, v_t, x^{t-1}) \cdot \chi_{ijt}(\theta_0, (\hat{\theta}_{it}^A, \theta^{-l,A}), (\hat{b}_{it}^A, v_t^{-l,A}), x^{t-1}) \quad (29)$$

$$+ \mathbb{E}^{\bar{\lambda}[\chi^{I;\beta}]|\theta_t^A, \theta^{-l,A}, v_{it}^A, v_t^{-l,A}, x^{t-1}; (\hat{\theta}_{it}^A, \hat{b}_{it}^A)} [W_{t+1}(\theta_0, \theta_{t+1}, b_{t+1}, x^t)].$$

The inequality in (29) follows from the fact that the matches that obtain when all agents follow truthful strategies from period t (included) onwards maximizes continuation weighted surplus.

Step 2. To complete the proof of part 1, it then suffices to show that each agent finds it optimal to participate in period 0 and select the membership status designed for his true vertical type. As in the case of exogenous processes, this follows from the fact that the matches sustained under truthful strategies satisfy monotonicity conditions analogous to those in Lemma 2. The arguments are similar to those in the proof of Theorem 1 and hence are omitted.

Part 2. Let $\chi_t^P(\theta, \omega)$ and $\chi_t^W(\theta, \omega)$ denote the state-contingent matches implemented in period $t \geq 1$, under the truthful equilibria of, respectively, the profit-maximizing and the welfare-maximizing auctions. Note that the arguments of these functions are the exogenous vertical types θ and the sequences of exogenous innovations $\omega \equiv (\omega_{ijs}^k)_{(i,j) \in N^A \times N^B, k=A,B}^{s=1, \dots, \infty}$ that, along with the matches implemented in previous periods generate the horizontal types ε . Because the match values are endogenous, such a representation favors the comparison of the matches sustained under the two auctions by making the “state variables” exogenous, thus eliminating the confusion that may originate from the fact that the histories of horizontal types need not coincide under the two auctions.³⁷

Similarly, let $S_{ijt}^{I;P}(\theta, \omega)$ and $S_{ijt}^{I;W}(\theta, \omega)$ denote the period- t indexes under the truthful equilibria of the profit-maximizing and the welfare-maximizing auctions, respectively. Because (θ, ω) are exogenous and time-invariant, they are dropped from all the functions χ^P , χ^W , $S^{I;P}$, and $S^{I;W}$ below.

First, observe that, because $\beta_l^{k,P}(\theta_l^k) \leq 1 = \beta_l^{k,W}(\theta_l^k)$, all $\theta_l^k \in \Theta_l^k$, $l \in N^k$, $k = A, B$, and because the evolution of the match values is time-autonomous and the horizontal types are nonnegative, for any $(i, j) \in N^A \times N^B$, any $t, \tau \geq 1$,

$$\sum_{s=1}^{t-1} \chi_{ijs}^W = \sum_{s=1}^{\tau-1} \chi_{ijs}^P \Rightarrow S_{ijt}^{I;W} \geq S_{ij\tau}^{I;P}. \quad (30)$$

³⁷Recall that, given the matches x , the horizontal types ε can be represented as deterministic functions of (x, ω) , as explained at the beginning of Section 6.

When none of the capacity constraints binds, in each period $t \geq 1$, the matches implemented under the equilibria of the profit-maximizing auctions (alternatively, the welfare-maximizing auctions) are all those for which the index $S_{ijt}^{I;P} \geq 0$ (alternatively, $S_{ijt}^{I;W} \geq 0$). The result that, for all $(i, j) \in N^A \times N^B$, all $t \geq 1$, $\chi_{ijt}^P = 1 \Rightarrow \chi_{ijt}^W = 1$ then follows from the fact that, for any $(i, j) \in N^A \times N^B$, $t \geq 1$, $S_{ijt}^{I;W} \geq S_{ijt}^{I;P}$. The last property, in turn, follows by induction. First observe that the property is necessarily true at $t = 1$, given (30) and the fact that, in period $t = 1$, the number of past interactions is necessarily the same under profit and welfare maximization. Now suppose the result holds for all $1 \leq s < t$. Note that any match for which $S_{ijt}^{I;P} \geq 0$ has been active at each preceding period $s < t$, both under profit maximization and under welfare maximization. The result then follows again from (30), which implies that $S_{ijt}^{I;W} \geq S_{ijt}^{I;P}$.

Now suppose $M = 1$ but $m_l^m \geq n^{-k}$, all $l \in N^k$, $k = A, B$, all $t \geq 1$. We want to show that, for all $t \geq 1$, $\sum_{(i,j) \in N^A \times N^B} \chi_{ijt}^W \geq \sum_{(i,j) \in N^A \times N^B} \chi_{ijt}^P$. Because $M = 1$, this means that, in any period in which matching is active under profit maximization, it is also active under welfare maximization.

First observe that, under the equilibria of the profit-maximizing auction, if at some period $t \geq 1$, $\chi_{ijt}^P = 0$, all $(i, j) \in N^A \times N^B$, then $\chi_{ijs}^P = 0$, all $s > t$, all $(i, j) \in N^A \times N^B$. The same property holds for χ^W . Next, observe that, if matching stops at period t under profit maximization (alternatively, welfare maximization), then $S_{ijt}^{I;P} < 0$ all $(i, j) \in N^A \times N^B$ (alternatively, $S_{ijt}^{I;W} < 0$ all $(i, j) \in N^A \times N^B$). Now suppose that, under profit maximization, matching is still active in period t (meaning, there exists $(i, j) \in N^A \times N^B$ such that $\chi_{ijt}^P = 1$). Then there are two cases. (1) Either $\sum_{s=1}^{t-1} \chi_{ijs}^W = \sum_{s=1}^{t-1} \chi_{ijs}^P$, all $(i, j) \in N^A \times N^B$, in which case (30) implies that $S_{ijt}^{I;W} \geq S_{ijt}^{I;P}$ for all $(i, j) \in N^A \times N^B$, and the result then obviously holds. Or, (2) there exists $(i, j) \in N^A \times N^B$ such that $\sum_{s=1}^{t-1} \chi_{ijs}^W < \sum_{s=1}^{t-1} \chi_{ijs}^P$. To see this, recall that the fact that matching is still active in period t under profit maximization implies it must have been active in each preceding period as well. That $\sum_{s=1}^{t-1} \chi_{ijs}^W < \sum_{s=1}^{t-1} \chi_{ijs}^P$ in turn implies there must exist $\tau < t$ such that $\sum_{s=1}^{\tau-1} \chi_{ijs}^P = \sum_{s=1}^{t-1} \chi_{ijs}^W$, and $\chi_{ij\tau}^P = 1$. That $\chi_{ij\tau}^P = 1$ in turn implies $S_{ij\tau}^{I;P} \geq 0$. That $\sum_{s=1}^{\tau-1} \chi_{ijs}^P = \sum_{s=1}^{t-1} \chi_{ijs}^W$ in turn implies that $S_{ijt}^{I;W} \geq S_{ij\tau}^{I;P}$, where the result follows again from (30). From the discussion above, that $S_{ijt}^{I;W} \geq S_{ij\tau}^{I;P}$ in turn implies that matching must be active in period t also under welfare maximization. This completes the proof of Part 4 and of the theorem. Q.E.D.

References

- [1] Akbarpour, M., Li, S., and S. Oveis Gharan (2016), “Thickness and Information in Dynamic Matching Markets,” mimeo Stanford University.
- [2] Allard, N., (2008), “Lobbying Is An Honorable Profession: The Right To Petition and The Competition To Be Right,” *Stanford Law Review*, Vol. 19(1), 23-68.
- [3] Anderson, R., Ashlagi, I., Gamarnik, D., and Y. Kanoria (2015), “Efficient Dynamic Barter Exchange,” mimeo, MIT.
- [4] Arnosti, N., Beck, M., and P. Milgrom (2016), “Adverse Selection and Auction Design for Internet Display Advertising,” *American Economic Review*, 106(10), 2852-66.
- [5] Asker, J., and E. Cantillon (2008), “Properties of Scoring Auctions,” *The RAND Journal of Economics*, 39, 69-85.
- [6] Athey, S. and G. Ellison (2011), “Position Auctions with Consumer Search,” *Quarterly Journal of Economics*, Vol. 126, 1213–1270.
- [7] Baccara, M., S. Lee, and L. Yariv (2015): “Optimal Dynamic Matching,” mimeo Caltech.
- [8] Baron, D. and D. Besanko (1984), “Regulation and Information in a Continuing Relationship,” *Information Economics and Policy*, 1, 447–470.
- [9] Belleflamme, P., and M. Peitz, (2017), “Platforms and Network Effects,” In Corchon, L. and M. Marini (Eds). *Handbook of Game Theory and Industrial Organization*. Edward Elgar. Cheltenham, UK, Northampton, MA, USA.
- [10] Bergemann, D. and A. Pavan (2015), “Introduction to Symposium on Dynamic Contracts and Mechanism Design,” *Journal of Economic Theory* 159, 679-701.
- [11] Bergemann, D. and J. Välimäki (2010), “The Dynamic Pivot Mechanism,” *Econometrica*, 78(2), 771–789.
- [12] Bergemann, D. and J. Välimäki (2017), “Dynamic Mechanism Design: An Introduction,” mimeo, Yale University.
- [13] Besanko, D. (1985), “Multiperiod Contracts Between Principal and Agent with Adverse Selection,” *Economics Letters*, 17, 33–37.
- [14] Bloch, F. and D. Cantala (2016), “Dynamic Allocation of Objects to Queuing Agents,” *American Economic Journal: Microeconomics*, forthcoming.
- [15] Board, S. (2007), “Selling Options,” *Journal of Economic Theory* 136, 324–340.
- [16] Board, S. (2009), “Monopoly Group Design with Peer Effects,” *Theoretical Economics*, Vol. 4, 89-125.

- [17] Borgers, T. (2015), “An Introduction to the Theory of Mechanism Design,” Oxford University Press, Oxford.
- [18] Borgers, T., Cox, I., Pesendorfer, M., and V. Petricek (2013), “Equilibrium Bids in Sponsored Search Auctions: Theory and Evidence,” *American Economic Journal: Microeconomics*, Vol. 5, 163-187.
- [19] Cabral L. (2011): “Dynamic Price Competition with Network Effects,” *Review of Economic Studies* 78 (1): 83-111.
- [20] Che, Y. K. (1993): “Design Competition through Multi-Dimensional Auctions,” *RAND Journal of Economics* Vol. 24, pp. 668-680.
- [21] Cremer, J. and R. P. McLean (1988), “Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions,” *Econometrica*, 56(6), 1247–1257.
- [22] Courty, P. and H. Li (2000), “Sequential Screening,” *Review of Economic Studies*, 67, 697–717.
- [23] Damiano, E. and H. Li (2007), “Price Discrimination and Efficient Matching,” *Economic Theory*, Vol. 30, 243-263.
- [24] Damiano, E. and R. Lam, (2005), “Stability in dynamic matching markets,” *Games and Economic Behavior*, 52(1):34–53
- [25] Dekel, E., Jackson, M.O. and A. Wolinsky, (2008) “Vote Buying: General Elections,” *Journal of Political Economy*, 116(2), 351–380.
- [26] Dizdar, D., and B. Moldovanu (2016). “On the Importance of Uniform Sharing Rules for Efficient Matching,” *Journal of Economic Theory*, 165, 106-123.
- [27] Doval, L. (2016), “A Theory of Stability in Dynamic Matching Markets,” mimeo Northwestern University.
- [28] Edelman, B., Ostrovsky, M., and M. Schwarz (2007), “Internet Advertising and the Generalized Second-Price Auction,” *American Economic Review*, Vol. 97, 242-259.
- [29] Esó, P. and B. Szentes (2007), “Optimal Information Disclosure in Auctions and the Handicap Auction,” *Review of Economic Studies*, 74(3), 705–731.
- [30] Fershtman, D. and A. Pavan (2017), “Pandora’s Auctions: Dynamic Matching with Unknown Preferences,” *American Economic Review P&P*, 107(5), 186-90.
- [31] Gershkov, A. and B. Moldovanu (2014), “Dynamic Allocation and Pricing: A Mechanism Design Approach,” MIT Press.
- [32] Gomes, R. (2014), “Optimal auction design in two-sided markets,” *The RAND Journal of Economics*, 45: 248–272.

- [33] Gomes, R. and A. Pavan (2016), “Many-to-Many Matching and Price Discrimination,” *Theoretical Economics*, 11(3), 1005–1052.
- [34] Gomes, R. and A. Pavan (2018), “Price Customization and Targeting in Platform Markets,” mimeo Northwestern University and Toulouse School of Economics.
- [35] Gomes, R. and A. Sweeney (2014), “Bayes–Nash Equilibria of the Generalized Second-price Auction,” *Games and Economic Behavior*, 86, 421–437.
- [36] Herbst, H., and B. Schickner (2016), “Dynamic Formation of Teams: When Does Waiting for Good Matches Pay Off?,” mimeo University of Bonn.
- [37] Hoppe, H. C., Moldovanu, B., and A. Sela (2009), “The Theory of Assortative Matching Based on Costly Signals,” *Review of Economic Studies*, 76, 253–281.
- [38] Hoppe, H. C., Moldovanu, B., and E. Ozdenoren (2011). “Coarse Matching with Incomplete Information,” *Economic Theory*, 47(1), 75–104.
- [39] Johnson, T., (2013), “Matching Through Position Auctions,” *Journal of Economic Theory*, 148, 1700–1713.
- [40] Jullien, B. and A. Pavan (2018), “Information Management and Pricing in Platform Markets,” *Review of Economic Studies*, forthcoming.
- [41] Kakade, S., Lobel, I., and H. Nazerzadeh (2011), “Optimal Dynamic Mechanism Design and the Virtual Pivot Mechanism,” *Operations Research*, 61, 837–854.
- [42] Kang, K. and H. You (2016), “Lobbyists as Matchmakers in the Market for Access,” mimeo Carnegie Mellon University.
- [43] Kurino, M. (2009) “Credibility, efficiency and stability: a theory of dynamic matching markets,” *Jena economic research papers, JENA*.
- [44] Leshno, J. (2015), “Dynamic Matching in Overloaded Waiting Lists,” mimeo Columbia University.
- [45] Mansour, Y., S. Muthukrishnan, and N. Nisan (2012) “Doubleclick Ad Exchange Auction,” mimeo, Tel Aviv University.
- [46] Moghaddam, M and S. Y. Nof (2016), “Best Matching: Theory and Applications,” Springer.
- [47] Pavan, A. (2017), “Dynamic Mechanism Design: Robustness and Endogenous Types,” *Advances in Economics and Econometrics*, Proceedings of the 11th World Congress of Econometric Society, Volume 1, 1-62.
- [48] Pavan, A., I. Segal, and J. Toikka (2014), “Dynamic Mechanism Design: A Myersonian Approach,” *Econometrica*, 82, 601–653.
- [49] Schummer, J., (2016), “Influencing Waiting Lists,” mimeo Northwestern University.

- [50] Thakral, N., (2015), “Matching with Stochastic Arrival: Theory and an Application to Public Housing,” mimeo Harvard University.
- [51] Ünver, U. (2010), “Dynamic Kidney Exchange,” *Review of Economic Studies*, 77, 372-414.
- [52] Varian, H. (2007), “Position Auctions,” *International Journal of Industrial Organization*, Vol. 25, pp. 1163- 1178.
- [53] Varian, H. and C. Harris (2014). “The VCG Auction in Theory and Practice,” *American Economic Review*, 104(5), 442-45.
- [54] Whittle, P. (1982): *Optimization Over Time*, vol. 1. Wiley, Chichester.